Inducing Uncertainty on Open-Weight Models for Test-Time Privacy in Image Recognition

Abstract

A key concern for AI safety remains understudied in the machine learning (ML) literature: how can we ensure users of ML models do not leverage predictions on incorrect personal data to harm others? This is particularly pertinent given the rise of open-weight models, where simply masking model outputs does not suffice to prevent adversaries from recovering harmful predictions. To address this threat, which we call *test-time privacy*, we induce maximal uncertainty on protected instances while preserving accuracy on all other instances. Our proposed algorithm uses a Pareto optimal objective that explicitly balances test-time privacy against utility. We also provide a certifiable approximation algorithm which achieves (ε, δ) guarantees without convexity assumptions. We then prove a tight bound that characterizes the privacy-utility tradeoff that our algorithms incur. Empirically, our method obtains at least $> 3\times$ stronger uncertainty than pretraining. Altogether, this framework provides a tool to guarantee protection to end users.

1 Introduction

Data privacy is increasingly important for large-scale machine learning (ML), where models are often trained on sensitive user instances (GDPR, 2016). Furthermore, open-weight image recognition models, where users have access to the model parameters and architecture, have proliferated (TorchVision, 2016; Google, 2023; Microsoft, 2024).

Yet, there has been little work done to address privacy threats to ML models due to incorrect personal data, especially data which are public such as images posted to public forums. Concretely, suppose a model provider trains an open-weight medical imaging model f which classifies skin images as harmless ailments like "Benign Keratosis" or serious diseases like "Melanoma" (Sun et al., 2016). Next, a health insurance company scrapes images from public forums to build risk profiles. Then, this health insurance company downloads the open-weight model f to automatically screen images for potential health liabilities. In particular, a person p posts a photo of a harmless birthmark to a public health forum to ask a question. During the upload, a compression error causes the image file z_p to become corrupted, severely distorting the birthmark. This results in a publicly available but incorrect image x_p; note that the original z_p is *never publicly available*. When the health insurance company feeds x_p, scraped from the public forum, into f, it confidently classifies x_p as "Melanoma." This erroneous classification is then automatically added to person p's risk profile, resulting in person p being unfairly denied coverage.[1] We call this threat model *test-time privacy* (TTP), and make this concrete in Fig. 1.[2] This privacy threat model is inspired by definitions of privacy which correspond to protecting a user from unauthorized interference or intrusion (Merriam-Webster, 2022). This differs from settings in privacy which mainly protect sensitive information.

---

[1] Recently, there has been significant progress in building effective image recognition models for skin disease, making this problem pertinent (Yang et al., 2018); (Liu et al., 2025).

[2] We provide additional test-time privacy examples beyond medical classifiers in Appx. D.

Figure 1: An adversary, like a health insurance company ( ), can query a pretrained model f ( ) and use its outputs to make harmful decisions. However, after running our algorithm, the new model f_U ( ) provides maximal uncertainty, protecting against such an adversary.

We discuss why existing solutions, like unlearning (Sekhari et al., 2021) or differential privacy (Dwork et al., 2006), do not suffice to solve this problem in Sec. 2. Furthermore, naive solutions like masking model outputs do not work for open-weight image recognition models, since the model parameters and architecture are available to the model user. An adversary could simply remove such a mask. To make this clear, we comprehensively detail our threat model as a security game in Sec. 3 and provide various motivating attacks in Appx. A. Therefore, we ask the following research question:

Can we ensure test-time privacy against adversaries with access to an open-weight model?

To do so, we argue it suffices to have uniform model outputs over the protected instances. That way, a data controller can only guess at the prediction. Thus, we revisit inducing maximal uncertainty over a dataset (Pereyra et al., 2017). Furthermore, we want to obtain high performance on all other instances as well. In particular, we answer our research question affirmatively, providing:

- A method to finetune a pretrained model with a Pareto optimal objective, rendering the model maximally uncertain over protected instances while preserving accuracy on others.

- Principled (ε, δ)-certified algorithms which approximate the Pareto objective, derived without assumptions of convexity.

- A theoretical analysis of the privacy-utility tradeoff that our algorithms incur, establishing a tight, non-vacuous bound.

- Empirical studies on image recognition models like ResNet50 (He et al., 2016) trained on datasets like CIFAR100 (Krizhevsky et al., 2009), observing that our algorithms maintain high uniformity on protected instances while guaranteeing excellent utility on the rest. In the appendix, we also provide results on a ViT (Dosovitskiy et al., 2021) pretrained on ImageNet (Deng et al., 2009), obtaining almost perfect uniformity on protected instances with a $< 1\%$ drop in accuracy on others.

Following the literature on privacy, we focus on protecting a subset of the training data. However, as detailed in Sec. 3, our setup and algorithms can also work for corrupted test instances. The code for our experiments is available for reproducibility at `https://anonymous.4open.science/r/ensuring_test_time_privacy-1C7B/README.md`.

2 Related Work

Data Privacy: Data leakage is a persistent danger for large information systems (Al-Rubaie & Chang, 2019). In the context of ML, data privacy is ubiquitous (Fredrikson et al., 2015); (Song et al., 2017); (Yeom et al., 2018). Approaches to privacy-preserving ML include differential privacy (DP) (Dwork et al., 2006); (Balle & Wang, 2018); (Amin et al., 2024), homomorphic encryption (Brakerski et al., 2014); (Lee et al., 2022); (Aono et al., 2017), and model obfuscation (Zhou et al., 2023). Notably, these methods protect against various privacy violations, like reconstruction attacks (Dinur & Nissim, 2003) due to failures of anonymization (Li et al., 2012). However, existing methods do not prevent confidently correct classification, and thus fail to protect against the attacks we consider in our setting. For example, if x_p is a corrupted medical image record, an adversary may not be able to use a DP model f to recover the record x_p exactly, but they can still produce a confident prediction of e.g. "Melanoma" to harm person p.

A dominant viewpoint in the privacy community is that a model f working as expected does not constitute a privacy violation, e.g. correctly predicting "Melanoma" for the corrupted medical image x_p, as it has learned something underlying about nature (Mcsherry, 2016); (Bun et al., 2021). Furthermore, the privacy leakage occurred when x_p became public (Kamath, 2020). This view misses the point: ML models are often trained and applied on freely available data. For example, training data could be scraped from the web or social media platforms. Subsets of this data can be obsolete, corrupted, or confidential. With such data as input, model f presents a clear and present danger for AI safety which differential privacy falls short of addressing, as the example in Sec. 1 shows. In parallel, as humans, we learn to not act upon certain kinds of knowledge. For example, when we read confidential documents or learn that previously obtained knowledge is incorrect, we are not allowed to share or act on this knowledge.

Unlearning: A related subfield is machine unlearning, which is inspired by the right to be forgotten (RTBF), mandating that ML model providers delete user data upon request (GDPR, 2016). In practice, model providers must remove user data and its effects from trained models and algorithms. Unlearning methods usually do so by approximating (and evaluating performance against) the model retrained from scratch without the protected user data (Sekhari et al., 2021); (Bourtoule et al., 2021); (Kurmanji et al., 2023).

However, while unlearning helps model providers comply with the RTBF, it cannot protect against attacks within our threat model. Specifically, recent unlearning research has established that data in the support of the training distribution will likely still be confidently predicted with the same prediction as before, even after using state-of-the-art algorithms (or even after applying exact retrain-from-scratch algorithms) to unlearn them (Zhao et al., 2024). That is, denoting the pretrained model as f and the unlearned model as f_u, for typical training instances, it holds that $f \approx f_u$.

To make clear why unlearning does not solve our problem, recall the example from Sec. 1 of a model f trained on skin images to predict disease. This time, to remove the corrupted medical image x_p from f, person p invokes the RTBF. Thus, the data controller for f unlearns x_p, yielding f_u. But, even after unlearning, any medical insurance company can still access the publicly available x_p and obtain $f_u(x_p)$. But, $f_u(x_p) \approx f(x_p)$, and thus the medical insurance company *incorrectly* labels person p as high risk for medical coverage. Thus, the unfair and dangerous scenario for person p remains. This holds similarly for unlearning methods which deal with corrupted or obsolete data, as they still do not aim to reduce confidence in the final prediction (Schoepf et al., 2025).

Differences from Unlearning: Importantly, what we propose is **not an unlearning algorithm**, which would need to be aligned with the goals of unlearning (and indistinguishability from retrain-from-scratch). Instead, we aim to address an entirely new threat scenario—test time privacy—that unlearning cannot solve, which we detail in Sec. 3. For example, indistinguishability from retrain is inconsequential in our threat model. Furthermore, we also consider corrupted test examples, unlike unlearning which focuses only on the training dataset. Finally, what may constitute a privacy violation in unlearning, e.g., revealing that an instance is in the protected set via a membership inference attack (Shokri et al., 2017) does *not* constitute a violation in our threat model. This holds similarly for other threat models; for example, reconstruction attacks which lead to recovery of x_p (Dinur & Nissim, 2003) or adversarial attacks which lead to misclassification of x_p (Goodfellow et al., 2015) are not violations in our threat model, as explained in Sec. 3.

Additional Related Works: We provide an additional related works section in Appx. B. Critically, we describe how differentially private methods like private aggregation of teacher ensembles (PATE) (Papernot et al., 2018) or label differential privacy (LabelDP) (Ghazi et al., 2021) differ from our setting, how label model inversion attacks (Zhu et al., 2019) relate to our threat model, why misclassification-based methods for unlearning (Cha et al., 2024) are suboptimal for addressing our threat model, and how our work relates to out-of-distribution and data abstention methods (Lee et al., 2018).

3 Test-Time Privacy Threat Model

Notation: Let $\mathcal{X} \subset \mathbb{R}^d$ be a sample space and let $\mathcal{Y} \subset \mathbb{R}^o$ be a label space. Denote $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ as the space of feature-label pairs. Let $\mathcal{Z}^n$ be the n-fold Cartesian product of $\mathcal{Z}$ such that a dataset $\mathcal{D} \subset \mathcal{Z}^n$ is a collection of n feature-label pairs. Then, the ith instance is denoted as $\mathcal{D}^{(i)}$ with its feature in $\mathcal{X}$ being $\mathcal{D}^{(i,\mathcal{X})}$ and label being $\mathcal{D}^{(i,\mathcal{Y})}$. A (randomized) learning algorithm is given by $\mathcal{A} : \mathcal{Z}^n \to \mathcal{W}$, where $\mathcal{W} \subset \mathbb{R}^z$ is a parameter space. Let the set of hypotheses parameterized with respect to this parameter space be $\mathcal{H}_{\mathcal{W}}$. Let $f_{\boldsymbol{w}} \in \mathcal{H}_{\mathcal{W}}$ be the hypothesis parameterized by $\boldsymbol{w} \in \mathcal{W}$, defined as $f_{\boldsymbol{w}} : \mathcal{X} \to \Delta_{|\mathcal{Y}|}$, where $\Delta_{|\mathcal{Y}|}$ is the probability simplex $\{p_1, ..., p_\gamma : p_i \geq 0, \sum_{i=1}^{|\mathcal{Y}|} p_i = 1\}$. When $\boldsymbol{A}$ is a matrix, $||\boldsymbol{A}||_2$ is the 2 operator norm. When $\boldsymbol{v}$ is a vector, $||\boldsymbol{v}||_2$ is the ℓ_2 norm. Furthermore, let $\lambda_{\min}(\boldsymbol{A})$ denote the minimum eigenvalue of $\boldsymbol{A}$. If we have an objective $f_A + f_B$, we denote its gradient evaluated at $\boldsymbol{w}$ as $\nabla_{\boldsymbol{w},A,B}$ and Hessian as $\boldsymbol{H}_{\boldsymbol{w},A,B}$. Finally, when for sets $\mathcal{S}, \mathcal{F} \subset \mathcal{N}$ and $\mathcal{R}$ and mechanisms $\mathcal{M}, \mathcal{M}' : \mathcal{N} \to \mathcal{R}$, we have $\Pr(\mathcal{M}(\mathcal{S}) \in \mathcal{R}) \leq e^\varepsilon \Pr(\mathcal{M}'(\mathcal{F}) \in \mathcal{R}) + \delta$ and $\Pr(\mathcal{M}'(\mathcal{S}) \in \mathcal{R}) \leq e^\varepsilon \Pr(\mathcal{M}(\mathcal{F}) \in \mathcal{R}) + \delta$, we will denote $\mathcal{M}(\mathcal{S}) \approx_{\varepsilon,\delta,\mathcal{R}} \mathcal{M}'(\mathcal{F})$. We provide a symbol table in Appx. K.

Following recent work (Salem et al., 2023), we define our threat model as an informal *test-time privacy (TTP) game*. In our setting, there are three core actors: a person p to protect, a benign model provider τ, and a TTP adversary ν. Firstly, a person p owns a true data instance $\boldsymbol{z}_p$ (e.g., the true birthmark image) which gets corrupted and released as instance $\boldsymbol{x}_p$ (e.g, the distorted birthmark image). Secondly, a model provider τ runs an algorithm $\mathcal{A}$ on a train set $\mathcal{D}$ to obtain initial model f. Then, the model provider is made aware of the corrupted $\boldsymbol{x}_p$ through, e.g., a direct petition by person p or a third-party entity. This is reasonable when person p is aware that their public information, like a corrupted medical image or incorrect credit report, can be used to harm them. Thus, τ re-releases a model $\hat{f}$ after running $\mathcal{G}$ on f and $\mathcal{D}$. Thirdly, a TTP adversary ν aims to use $\hat{f}$ to get a confident, harmful prediction m (e.g., Melanoma) on the public $\boldsymbol{x}_p$ to use as an automated warrant for some downstream task (e.g., rejecting insurance applications).

Given this scenario, the adversary ν wins if they can leverage open-weight access to $\hat{f}$ to *confidently* obtain the harmful prediction m. The provider τ wins if $\hat{f}$ leaves the adversary ν uncertain about m. To present our threat model in the clearest fashion, we make two key assumptions:

Assumption 3.1 (Open-Weight Access). *The adversary ν has full access to $\hat{f}$'s architecture and weights.*

Assumption 3.2 (τ-Limited Knowledge). *The provider τ has access to $\boldsymbol{x}_p$ but does not know the specific harmful label m the adversary seeks, nor do they know the adversary ν.*

Discussion: Firstly, Asm. 3.1 provides clear scope for our problem setting. Otherwise, the threat model becomes significantly more complicated: the provider τ can simply release $\hat{f}$ with a mask on $\boldsymbol{x}_p$ that provides uniform outputs, and the adversary ν could not remove it easily and would have to employ model stealing (Carlini et al., 2024) to do so. In this case, ν would have access to an inexact model $\hat{g}$ which aims to approximate $\hat{f}$ rather than $\hat{f}$ itself. Instead, to present our threat model and solutions in the clearest manner, we assume a stronger setting where ν has full access to $\hat{f}$. Secondly, Asm. 3.2 is reasonable since the model provider has no way of reliably knowing what downstream task ν is aiming for or what precisely their released model will be used for. Asm. 3.2 also rules out misclassification (Cha et al., 2024) or relabeling (Angelopoulos et al., 2025) of $\boldsymbol{x}_p$. Suppose that, following the example provided in Sec. 1, the model provider relabels the instance $\boldsymbol{x}_p$ as "Benign". But, the model provider τ does not know whether ν is a medical insurance company aiming to obtain a prediction of "Melanoma" to reject coverage or a defense attorney in a criminal case against a doctor aiming to obtain a prediction of "Benign" to clear a doctor of accusations of medical malpractice, thus failing to win the game.

Under Assumptions 3.1 and 3.2, for a confidence threshold $\rho \in (\frac{1}{|\mathcal{Y}|}, 1)$ and harmful label $m \in [|\mathcal{Y}|]$, the adversary ν wins if:

$$\hat{f}(\boldsymbol{x}_p)_m \geq \rho.$$

While the adversary ν knows the harmful label m, the model provider does not under Asm. 3.2. Thus, the provider seeks to reduce the worst-case confidence, i.e. solve:

$$\min_{f(\boldsymbol{x}_p) \in \Delta_{|\mathcal{Y}|}} \max_{j \in [|\mathcal{Y}|]} f(\boldsymbol{x}_p)_j$$

This quantity is minimized over the probability simplex by the uniform distribution, i.e. setting:

$$\hat{f}(\boldsymbol{x}_p)_j := \frac{1}{|\mathcal{Y}|} \quad \forall j \in [|\mathcal{Y}|].$$

That is, the provider's optimal defense is to have $\hat{f}$ output maximal uncertainty (i.e., a uniform distribution) on $\boldsymbol{x}_p$ while maintaining accuracy on all other data. Due to Asm. 3.1, this should be done in a way which perturbs the model weights, which motivates our approach in Sec. 4. This threat model generalizes to a collection of $\boldsymbol{x}_p$, which we denote as the "protected set" $\mathcal{D}_p$. For notational convenience, we write $\mathcal{D}_p$ as the set of protected instances supplied to $\mathcal{G}$. These instances may or may not belong to the original training set $\mathcal{D}$: all guarantees provided in Sec. 4 and Sec. 5 follow the same way if $\mathcal{D}_p$ contains test instances as well.

Importantly, we consider a *malicious or negligent* adversary which, even if they know the data is corrupted, still tries to obtain the harmful prediction m using the native input $\boldsymbol{x}_p$ to use as an automated warrant to make decisions (e.g., denying insurance coverage). This rules out any approaches based on the model provider τ detecting or flagging corrupted data (Yang & Weng, 2025).

Next, our work is motivated by *prevention* of harm rather than *correction*: we do not aim to repair the public record, recover the true instance, or prevent harm from previously released models. Instead, we focus on preventing the current release $\hat{f}$ from serving as a highly confident automated warrant to harm person p.

Crucially, TTP is *not a robustness failure* in the usual sense of requiring correct classification under perturbations. Specifically, the model f classifying $\boldsymbol{x}_p$ as, e.g., Melanoma instead of Benign is not a failure of robustness. $\boldsymbol{x}_p$ is *incorrect*: the semantic content of $\boldsymbol{x}_p$ is significantly different than that of $\boldsymbol{z}_p$ even though they are from the same distribution. The model f is thus not necessarily learning the wrong decision rule by classifying $\boldsymbol{x}_p$ as, e.g., Melanoma. Given this, standard robustness (Madry et al., 2018) or out-of-distribution (Hendrycks et al., 2019; Lee et al., 2018) methods do not directly target our threat model. Similarly, this is not a failure of calibrating (Guo et al., 2017) the model either; the harmful classification is due to false attributes in $\boldsymbol{x}_p$, rather than miscalibration on $\boldsymbol{z}_p$.

Rather than a robustness failure, TTP is a *privacy problem*. The corrupted $\boldsymbol{x}_p$ instance is incorrectly publicly available and the use of this instance $\boldsymbol{x}_p$ instead of $\boldsymbol{z}_p$ violates person p's privacy as discussed in Sec. 2. This is also why a simple data integrity check, like checking the hash of $\boldsymbol{z}_p$, does not solve the problem, since $\boldsymbol{z}_p$ differs from $\boldsymbol{x}_p$. Notably, person p's attribution to $\boldsymbol{x}_p$ is part of the metadata of $\boldsymbol{x}_p$, and cannot be obscured without altering public information about $\boldsymbol{x}_p$, which is out of scope. Thus, while necessary for our threat model, the core privacy concern is not identifying person p with $\boldsymbol{x}_p$, and thus not the concern we aim to protect against. Due to this, membership inference attacks (Shokri et al., 2017) and reconstruction attacks (Haim et al., 2022) are not the focus of our work. We discuss tradeoffs with membership inference in Appx. E. Additionally, while person p may have access to $\boldsymbol{z}_p$, and even can share it privately or publicly, the incorrect instance $\boldsymbol{x}_p$ can still be used by the adversary to cause harm; publishing $\boldsymbol{z}_p$ will not remove $\boldsymbol{x}_p$'s existence from the public record. Furthermore, $\mathcal{D}_p$ is supplied to the provider τ through an authenticated or externally validated process; designing such a process and preventing malicious petitions are orthogonal governance questions outside our threat model's scope, since we assume that person p and the model provider are benign.

Finally, in Appx. A, we present some simple TTP attacks on open-weight image classifiers to further motivate our threat model. Resistance to such attacks is an empirical strengthening of the basic TTP objective, which we evaluate separately in Appx. J.7. We also provide further motivating examples of TTP in Appx. D.

4 Approaches and Algorithms

As detailed in Sec. 3, it suffices to have the model output a uniform distribution over the protected set, rendering the model maximally uncertain. Then, an adversary cannot obtain a high confidence prediction on $\boldsymbol{x}_p$ to use as a warrant. We also would like to preserve retain set accuracy; to that end, we present an algorithm that finetunes the pretrained model with a Pareto optimal objective. To make this algorithm concrete, we define a *uniform learner*, which we prove to exist in many common hypothesis classes. Then, we use this concept in order to construct a Pareto objective.

4.1 The Finetuning Pareto Learner

For a dataset $\mathcal{D} \subset \mathcal{Z}^n$, we denote the pretrained model as $\mathcal{A}(\mathcal{D})$. Then, to make $\mathcal{A}(\mathcal{D})$ uniform over the protected set, we introduce our core concept of a *uniform learner*:

Definition 4.1 (Uniform learner). *Suppose we have a (randomized) learning algorithm $\mathcal{K} : \mathcal{Z}^n \to \mathcal{W}$ that, given $\mathcal{D} \subset \mathcal{Z}^n$, yields the parameter $\mathcal{K}(D) = \boldsymbol{w}_{\mathcal{D}}$. We say $\mathcal{K}$ is a uniform learner if $\forall \mathcal{D} \in \mathcal{Z}^n$, $\boldsymbol{w}_{\mathcal{D}}$ parametrizes $f_{\boldsymbol{w}_{\mathcal{D}}} \in \mathcal{H}_{\mathcal{W}}$ and satisfies:*

$$||f_{\boldsymbol{w}_{\mathcal{D}}} - (\underbrace{\frac{1}{|\mathcal{Y}|}, ..., \frac{1}{|\mathcal{Y}|}}_{|\mathcal{Y}| \; times})||_{\infty,\mathcal{X}} = 0. \tag{1}$$

That is, $\mathcal{K}$ is a uniform learner if its parameterized outputs yield the uniform distribution $U[0, |\mathcal{Y}|]$ for all inputs across all datasets. We define this as a learning algorithm for full generality to handle e.g. neural networks with nonlinear transformations in their last layer. Furthermore, $\mathcal{K}$ exists in many common hypothesis spaces, including those of multilayer perceptrons (MLP), residual networks (ResNets) (He et al., 2016), and transformers (Vaswani et al., 2017); this is proved in Appx. H.2.

Next, we assume $\mathcal{A}$ and $\mathcal{K}$ are obtained through empirical risk minimization (ERM) to a local or global minima. That is, let $\mathcal{A}(\mathcal{D}) = \mathrm{argmin}_{\boldsymbol{w} \in \mathcal{W}} \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D})$ and $\mathcal{K}(\mathcal{D}) = \mathrm{argmin}_{\boldsymbol{w} \in \mathcal{W}} \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D})$, where $\mathcal{L}_{\mathcal{A}}$ penalizes incorrect classification and $\mathcal{L}_{\mathcal{K}}$ penalizes a lack of uniformity in model outputs. One choice of $\mathcal{L}_{\mathcal{K}}$ is the KL divergence (Kullback & Leibler, 1951) between the softmax outputs and the uniform distribution. This loss has been previously used to penalize highly confident classifier predictions (Pereyra et al., 2017); we thus adapt this loss, which can be completely minimized over $\mathcal{W}$ by Prop. H.10, to our setting.

Critically, we seek the optimal tradeoff between uniformity over the protected set and utility over the retain set. That is, we should produce a learner that is Pareto optimal with respect to $\mathcal{L}_{\mathcal{K}}$ and $\mathcal{L}_{\mathcal{A}}$. In particular, if we take $\mathcal{M}_\theta(\mathcal{D}) = \mathrm{argmin}_{\boldsymbol{w} \in \mathcal{W}} \theta \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1 - \theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)$, we obtain a locally (globally) Pareto optimal learner, in the case that the local (global) minima is achieved. We prove this and provide a definition of Pareto optimality in Appx. H.3. Importantly, one can adjust θ to vary over many Pareto optimal solutions, yielding different tradeoffs between uniformity and utility. This yields Alg. 1, in which we finetune a pretrained model by using it as initialization for $\mathcal{M}_\theta(\mathcal{D})$.

While Alg. 1 is conceptually straightforward, this is an advantage of our approach and not a weakness for deployment in real-world settings as well as its reproducibility as a key baseline in future solutions to TTP. Alg. 1 is similar in spirit to algorithms from unlearning which aim to finetune a pretrained model over a subset of the training instances (Graves et al., 2021; Kurmanji et al., 2023). However, as detailed in Sec. 2, our threat model and resulting algorithm are different than those in unlearning. The approach in Alg. 1 is also similar to that of Lee et al. (2018), but differs in our objective formulation and also in that we do not assume that $\mathcal{D}_p$ is drawn from a different distribution than $\mathcal{D}_r$. Additionally, Alg. 1 is different from simply incorporating random labels over $\mathcal{D}_p$ and retraining from scratch. As demonstrated in Sec. 6 and Appx. J.12, approaches based on incorporating random labels during pretraining do not work as well as Alg. 1.

4.2 The Certified Pareto Learner

While Alg. 1 provably guarantees an optimal tradeoff, so long as its objective is minimized, we would also like to make it *certified*, obtaining a certificate that a third party can inspect to verify test-time privacy. Thus, to design a certified approximation algorithm, we take inspiration from certified unlearning (Zhang et al., 2024), which aims to add a small amount of structured noise such that the pretrained model becomes indistinguishable from the retrained model. In our setting, we would like to make the pretrained model indistinguishable from the solution to the Pareto objective. To do so, we define a new notion of (ε, δ)-indistinguishability and use this definition to design a certifiable algorithm, with results in Sec. 6.

Firstly, to motivate our definition, recall the definition of differential privacy (Dwork et al., 2006):

Definition 4.2 ((ε, δ)-differential privacy)**.** *Suppose we have privacy budgets $\varepsilon \in (0, 1)$ and $\delta > 0$. A randomized algorithm $\mathcal{M} : \mathcal{Z}^n \to \mathcal{W}$ satisfies (ϵ, δ)-differential privacy if $\forall \mathcal{T} \subset \mathcal{W}$ and $\forall \mathcal{D}, \mathcal{D}' \in \mathcal{Z}^n$ s.t. $||\mathcal{D} - \mathcal{D}'||_1 \leq 1$:*

$$\mathcal{M}(\mathcal{D}) \approx_{\varepsilon, \delta, \mathcal{T}} \mathcal{M}(\mathcal{D}'). \tag{2}$$

This guarantees that the algorithm $\mathcal{M}$ applied on a dataset is statistically indistinguishable from the same algorithm applied on all datasets different by one instance. One can leverage this definition to formalize certified unlearning (Sekhari et al., 2021):

Definition 4.3 ((ε, δ)-certified unlearning)**.** *Suppose we have privacy budgets $\varepsilon \in (0, 1)$ and $\delta > 0$. Consider $\mathcal{D} \subset \mathcal{Z}^n$ and let $\mathcal{D}_p \subset \mathcal{D}$ be the protected set to be unlearned, and $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_p$ be the retain set. $\mathcal{U} : \mathcal{Z}^n \times \mathcal{Z}^n \times \mathcal{W} \to \mathcal{W}$ is an (ε, δ)-certified unlearning algorithm if $\forall \mathcal{T} \subset \mathcal{W}$, we have:*

$$\mathcal{U}(\mathcal{D}, \mathcal{D}_p, \mathcal{A}(\mathcal{D})) \approx_{\varepsilon, \delta, \mathcal{T}} \mathcal{A}(\mathcal{D}_r). \tag{3}$$

This formalizes making $\mathcal{A}(\mathcal{D})$ indistinguishable from $\mathcal{A}(\mathcal{D}_r)$. In light of Def. 4.3, we seek to make $\mathcal{A}(\mathcal{D})$ indistinguishable from $\mathcal{M}_\theta(\mathcal{D})$. Thus, we provide the following new definition:

Definition 4.4 (($\varepsilon, \delta, \theta$)-certified Pareto learner)**.** *Suppose we have privacy budgets $\varepsilon \in (0, 1)$ and $\delta > 0$ with $\mathcal{D} \subset \mathcal{Z}^n$. Let $\mathcal{D}_p \subset \mathcal{D}$ be the protected set and let $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_p$ be the retain set. Suppose we have $\theta \in (0, 1)$ and $\mathcal{M}_\theta(\mathcal{D}) = argmin_{\boldsymbol{w} \in \mathcal{W}} \theta \mathcal{L}_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p) + (1 - \theta)\mathcal{L}_\mathcal{A}(\boldsymbol{w}, \mathcal{D}_r)$, where $\mathcal{K}$ is a uniform learner and $\mathcal{A}$ is a learning algorithm both obtained through ERM. An algorithm $\mathcal{G} : \mathcal{Z}^n \times \mathcal{Z}^n \times \mathcal{W} \to \mathcal{W}$ is a $(\varepsilon, \delta, \theta)$-certified Pareto learner if $\forall \mathcal{T} \subset \mathcal{W}$:*

$$\mathcal{G}(\mathcal{D}, \mathcal{D}_p, \mathcal{A}(\mathcal{D})) \approx_{\varepsilon, \delta, \mathcal{T}} \mathcal{M}_\theta(\mathcal{D}). \tag{4}$$

Discussion: Qualitatively, the conditions in Def. 4.4 mean that the model obtained by algorithm $\mathcal{G}$ is statistically indistinguishable from a model that is Pareto optimal between utility over the retain set and uniformity over the protected set. Here, we consider the classical setting of $\varepsilon \in (0, 1)$.[3] In our definition, ε characterizes the indistinguishability, δ the catastrophic failure probability, and θ the strength of the induced uniformity, i.e. θ adds and does not diminish interpretability. Next, note that satisfying Def. 4.3 and Def. 4.4 together is not possible for protected sets which overlap; thus, a model provider should adopt whichever approach corresponds to their threat model. Finally, the TTP definition mirrors Def. 4.3 and Def. 4.2, which also compare to a reference model. The implicit assumption with all of these definitions is that the reference model performs well for the task; thus, this is not a limitation of Def. 4.4.

One way we can design a first algorithm which satisfies Def. 4.4 is by taking a Newton step towards the Pareto model and applying structured Gaussian noise; this yields Alg. 2, which is certifiable as proved in

---

[3]Notably, Balle & Wang (2018) provide a way to achieve (ε, δ)-indistinguishability for $\varepsilon > 1$, and their technique can be adapted without loss of generality to our setting.

Algorithm 1 $\mathcal{M}_\theta$ Finetuning

Require: Dataset $\mathcal{D}$; protected set $\mathcal{D}_p$; pretrained model $\boldsymbol{w}^* = \mathcal{A}(\mathcal{D})$; tradeoff θ; e epochs.
Initialize $\mathcal{M}_\theta(\mathcal{D})$ with $\boldsymbol{w}^*$, run SGD for e epochs to yield $\boldsymbol{w}^-$.
return $\boldsymbol{w}^-$.

Algorithm 2 $(\epsilon, \delta, \theta)$-Certified Uniformity with Exact Inverse Hessian

---

Require: Dataset $\mathcal{D}$; protected set $\mathcal{D}_p$; pretrained model $\boldsymbol{w}^* = \mathcal{A}(\mathcal{D})$; ε, δ; θ; λ; C.
$\tilde{\boldsymbol{w}} \leftarrow \boldsymbol{w}^* - (\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda \boldsymbol{I})^{-1} \nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$.   // Derived in Appx. F.
Compute Δ as the bound in Eq. (F.18).
$\sigma = \frac{\Delta}{\epsilon} \sqrt{2 \ln(1.25/\delta)}$.
$\boldsymbol{w}^- \leftarrow \tilde{\boldsymbol{w}} + Y$, $Y \sim \mathcal{N}(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$.
return $\boldsymbol{w}^-$.

Appx. F. Using local convex approximation (Nocedal & Wright, 1999), in which we add a regularization term to the objective of the Pareto learner, we design Alg. 2 without any assumptions of convexity on the component loss functions.

In addition, Alg. 2 requires inverting a Hessian, which is computationally infeasible for practical neural networks e.g. ResNets, even after employing conjugate gradient methods (Nocedal & Wright, 1999) and Hessian vector product techniques (Pearlmutter, 1994). To resolve this issue, we also a proof-of-concept Alg. 3 in Appx. F, which computes an efficient estimator for the inverse Hessian (Agarwal et al., 2016). Furthermore, this algorithm does not assume convergence to a local minima for $\mathcal{A}(\mathcal{D})$, handling e.g. early stopping. While Alg. 2 has more hyperparameters than Alg. 1, they offer a certificate which can be used to verify use of our method by a third party; we present ways to reduce hyperparameters in Appx. G. Thus, Def. 4.4 and its associated algorithms Alg. 2 and Alg. 3 provide strong conceptual and theoretical basis for TTP solutions.

5 Theoretical Analysis

In what follows, we aim to analyze various properties of Alg. 1 and Alg. 2 to understand how to appropriately choose θ and the privacy-utility tradeoffs these algorithms incur. To clarify the notation used in this section, we include a symbol table in Appx. K. Firstly, we seek to understand how we can choose θ to guarantee uniformity over the protected set. To do so, we provide a constraint to be satisfied to ensure uniformity. Then, we provide an appropriate lower bound on θ to ensure the constraint is satisfied. In doing so, one obtains a bound on the privacy of our algorithm. We next want to obtain a bound on the utility of our algorithm. To that end, we upper bound the difference between the retain loss of the locally optimal learned model $\mathcal{A}(\mathcal{D}_r)$ and the locally optimal solution to the Pareto objective $\mathcal{M}_\theta(\mathcal{D})$. We obtain a tight, non-vacuous bound with respect to θ and characterize it asymptotically. Incorporating the two bounds provides a concrete characterization of the privacy-utility tradeoffs that occur when our algorithms are used.

In particular, across all algorithms, we make the pretrained model indistinguishable from

$$\mathcal{M}_\theta(\mathcal{D}) = \arg \min_{||\boldsymbol{w}||_2 \leq C, \boldsymbol{w} \in \mathcal{W}} \theta \mathcal{L}_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p) + (1 - \theta)\mathcal{L}_\mathcal{A}(\boldsymbol{w}, \mathcal{D}_r) + \frac{\lambda}{2}||\boldsymbol{w}||_2^2$$

$$= \arg \min_{||\boldsymbol{w}||_2 \leq C, \boldsymbol{w} \in \mathcal{W}} \theta \sum_{i=1}^{|\mathcal{D}_p|} \ell_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p^{(i)}) + (1 - \theta) \sum_{j=1}^{|\mathcal{D}_r|} \ell_\mathcal{A}(\boldsymbol{w}, \mathcal{D}_r^{(j)}) + \frac{\lambda}{2}||\boldsymbol{w}||_2^2,$$

where $\ell_\mathcal{K}$ and $\ell_\mathcal{A}$ are component loss functions corresponding to individual data instances in the protected and retain sets, respectively. Note that λ is present either as weight decay in the Pareto learning in Alg. 1 or as part of the local convex approximation in Alg. 2. Furthermore, note that the objective is constrained by $||\boldsymbol{w}||_2 \leq C$; we use this as a part of our local convex approximation when deriving Alg. 2; it is however unnecessary for Alg. 1. Similarly, unlearning methods assume this either implicitly or explicitly (Zhang et al., 2024). One can use projected gradient descent (Nocedal & Wright, 1999) during pretraining to satisfy this constraint.

Note that Alg. 1 has $\lambda \approx 0$. For Alg. 2, by Lemma H.9, our models $f_{\boldsymbol{w}^-}$ and $f_{\mathcal{M}_\theta(\mathcal{D})}$ have approximately the same outputs over $\mathcal{D}_p$, where $\boldsymbol{w}^-$ are the weights after applying one of our TTP algorithms. Hence, for any of our algorithms, to ensure indistinguishability from uniformity over $\mathcal{D}_p$, it suffices to ensure that $\mathcal{M}_\theta$ satisfies the following constraint:

$$||f_{\mathcal{M}_\theta(\mathcal{D})}(\mathcal{D}_p) - U[0, |\mathcal{Y}|]||_\infty \leq \gamma. \tag{5}$$

We then have the following bound on Eq. (5) with respect to θ, the proof of which is in Appx. H.9:

Proposition 5.1. *Let $\mathcal{M}_\theta(\mathcal{D})$ be the global solution to the Pareto objective. Choose, as surrogate losses, $\ell_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p^{(i)}) = D_{KL}(f_{\boldsymbol{w}}(\mathcal{D}_p^{(i)})||U[0, |\mathcal{Y}|])$, the KL divergence between the model outputs over the protected set and the uniform distribution, and $\ell_\mathcal{A}(\boldsymbol{w}, \mathcal{D}_r^{(j)}) = \mathbb{H}_{CE}(\mathcal{D}_r^{(j,\mathcal{Y})}, f_{\boldsymbol{w}}(\mathcal{D}_r^{(j,\mathcal{X})}))$, the cross entropy between model predictions and labels over the retain set. Then, $||f_{\mathcal{M}_\theta(\mathcal{D})}(\mathcal{D}_p) - U[0, |\mathcal{Y}|]||_\infty \leq \sqrt{2(\frac{1-\theta}{\theta}|\mathcal{D}_r| \ln |\mathcal{Y}|)}$.*

Proof Sketch: By using Prop. H.10 and the fact that $\mathcal{M}_\theta(\mathcal{D})$ is a global minimizer, we can yield a bound on $\mathcal{L}_\mathcal{K}$. Then, standard inequalities yield our result.

Then, we can choose θ as follows to guarantee Eq. (5), the proof of which is in Appx. H.10:

Corollary 5.2. *For a $\gamma > 0$, choosing $\theta \geq \frac{2|\mathcal{D}_r| \ln |\mathcal{Y}|}{\gamma^2 + 2|\mathcal{D}_r| \ln |\mathcal{Y}|}$ guarantees that Eq. (5) holds.*

Discussion: Note that Cor. 5.2 is well-defined in that $\theta \in (0, 1)$ for any choice of $|\mathcal{D}_r|, |\mathcal{Y}|$ and ε. Furthermore, Cor. 5.2 restricts $\mathcal{M}_\theta(\mathcal{D})$ to a subset of Pareto optimal solutions, but this does not render it no longer Pareto optimal; thus, our formulation as in Appx. F still holds in its entirety. Importantly, this is a sufficient but not necessary condition to satisfy Eq. (5).

Similarly, by Lemma H.9, we can study the effect of θ in $\mathcal{M}_\theta(\mathcal{D})$ on the (empirical) retain error on $\mathcal{D}_r$, after our algorithms are applied. To provide this bound, we require two key assumptions:

Assumption 5.3. *The gradients of $\ell_\mathcal{K}$ and $\ell_\mathcal{A}$ are Lipschitz in $\boldsymbol{w}$ with constants $\frac{P_\mathcal{K}}{|\mathcal{D}_p|}$ and $\frac{P_\mathcal{A}}{|\mathcal{D}_r|}$.*

Assumption 5.4. *The Hessians of $\ell_\mathcal{K}$ and $\ell_\mathcal{A}$ are Lipschitz in $\boldsymbol{w}$ with constants $\frac{F_\mathcal{K}}{|\mathcal{D}_p|}$ and $\frac{F_\mathcal{A}}{|\mathcal{D}_r|}$.*

Discussion: Note that these assumptions are only used to prove Thm. 5.5 and in Appx. F; we do not require them to prove all previously mentioned theorems. These assumptions, similar to those studied by Zhang et al. (2024), are less restrictive than those typically studied in certified unlearning Sekhari et al. (2021); importantly, we do not assume (strong) convexity of the losses.

We then present a tight, non-vacuous bound on the retain error after applying our algorithms:

Theorem 5.5. *Suppose Assumptions 5.3 and 5.4 hold, and let $P_\mathcal{K}, P_\mathcal{K}, F_\mathcal{K}, F_\mathcal{A}$ be as defined in Assumptions 5.3 and 5.4. Let $\alpha^* := \mathcal{L}_\mathcal{A}(\mathcal{A}(\mathcal{D}_r), \mathcal{D}_r)$ be the locally optimal (empirical) retain loss, achieved by $\mathcal{M}_\theta(\mathcal{D})$ when $\theta = 0$. Let $\alpha(\theta) := \mathcal{L}_\mathcal{A}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r)$ be the locally optimal retain loss obtained by $\mathcal{M}_\theta(\mathcal{D})$ when $\theta \in (0, 1)$. Suppose all weights used throughout are bounded by $||\boldsymbol{w}||_2 \leq C$. Then, for λ sufficiently large:*

$$|\alpha^* - \alpha(\theta)| \leq \mathcal{O}(\lambda C^2 \theta + C^2 \theta^2). \tag{6}$$

Proof Sketch: We subtract the first order conditions, by definition of α^* and $\alpha(\theta)$, to get an expression with respect to the gradients; plugging in an equivalent path integral expression and applying Lemma H.2 yields our desired result, with a full proof (including the full bound) in Appx. H.11.

Discussion: Three key hyperparameters should be kept small to ensure high retain accuracy: the ℓ_2 regularization coefficient λ, the max model weight magnitude C, and the Pareto frontier hyperparameter θ. In particular, large regularization coefficients take the model off the Pareto frontier. However, smaller or sparser weights are preferred, since the bound grows quadratically in C.

In addition, that when $\theta = 0$, the bound simplifies to 0, indicating that it is tight near 0. We demonstrate that it is tight near 1 in Appx. J. Importantly, Theorem 5.5 applies as stated to Alg. 2, since λ is large. However, for Alg. 1, λ in our experiments is not large enough. Still, as demonstrated in Tab. J.19, this bound holds

(a) Accuracy on retain set for baselines as well as Alg. 1 and Alg. 2 with $\theta = 0.75$.

(b) Confidence distance on protected set for baselines as well as Alg. 1 and Alg. 2 with $\theta = 0.75$.

Figure 2: Across datasets, observe a significant drop in confidence distance, where lower is better, for both our algorithms. We also observe that both algorithms provide strong accuracy on the retain set. We observe similar behavior for the test set in Appx. J, while the baselines are inconsistent.

empirically for Alg. 1. This bound thus provides a heuristic approximation for the behavior of Alg. 1, and a concrete result for Alg. 2. Furthermore, we can obtain a more concise bound with simpler techniques, but such a bound is vacuous and does not incorporate information about θ; we elaborate on this in Appx. H.11.

6 Empirical Analysis

Below, we provide empirical results for Alg. 1 and Alg. 2. We firstly discuss our experimental setup, baselines, and define our uniformity metric. Next, we provide our core results across Alg. 1, Alg. 2, and our baselines for several architectures and benchmarks. We also comment on the Pareto frontier of Alg. 1 and Alg. 2, providing additional insight into the structure of our problem.

Setup and Baselines. Our primary results on Alg. 1 are for ResNet50 (He et al., 2016) trained on SVHN, CIFAR10, and CIFAR100. We also provide results for logistic regression on MNIST to evaluate Alg. 2. We then include additional experiments with more complicated datasets and models, such as ViT (Dosovitskiy et al., 2021), TinyImageNet (Le & Yang, 2015), and ImageNet (Deng et al., 2009), in Tab. J.3. We compare results with the pretrained model and the model retrained without the protected set, which constitutes exact unlearning (Bourtoule et al., 2021). We also compare our methods to LabelDP (Ghazi et al., 2021) and a synthetic baseline that assigns random labels to instances neighboring the protected set. Across methods, we compare retain accuracy, test accuracy, and protected uniformity. We provide more details and the rationale for our baselines in Appx. I.

Providing a Uniformity Metric: We require a metric to compare uniformity over the protected set in an interpretable manner. Thus, we define the "confidence distance" as $f(\boldsymbol{x})_t - \frac{1}{|\mathcal{Y}|}$ for $\boldsymbol{x} \in \mathcal{D}_p$, where $f(\boldsymbol{x})_t$ is the max confidence score. In our experiments, we use this as the primary metric for uniformity, reporting the average confidence distance over the protected set. We discuss why this is reasonable in Appx. C and compare it to alternative metrics in Appx. J.

Overall Results: The results for Alg. 1 are presented in Fig. 2, in which we were able to achieve a $> 3\times$ decrease in confidence distance with only a 0.1% and 0.4% decrease in retain and test accuracy, respectively, for a ResNet50 pretrained on SVHN. For MNIST, CIFAR10, and CIFAR100, retain and test set accuracies incur a stronger drop. Still, they remain high, while protected confidence distance is significantly reduced. Results for the test set are deferred to Appx. J. We additionally find that the synthetic baseline can induce uniformity well for SVHN, but can either fail to induce uncertainty entirely (CIFAR10) or induce uncertainty at great cost to retain and test accuracy (CIFAR100). We observe similar behavior for TinyImageNet in Appx. J.3. This holds similarly for LabelDP, which furthermore undesirably reduces the confidence distance on retain and test sets, while our method does not, as demonstrated in Tab. J.4. Furthermore, our observations

(a) Retain Accuracy vs. θ.

(b) Confidence Distance vs. θ.

Figure 3: From Fig. 3a, we observe that for simple datasets, the retain accuracy decreases smoothly. However, for larger datasets like CIFAR10 and CIFAR100 as one passes $\theta \approx 0.75$, retain accuracy drops significantly. This motivates our choice of $\theta = 0.75$ used throughout our experiments. In Fig. 3b we observe that the confidence distance decreases roughly linearly as θ increases.

coincide with Zhao et al. (2024), observing that unlearned models still produce confident predictions on deleted instances.

Furthermore, as illustrated by Fig. 2, we find that Alg. 2 also induces uniformity well, while marginally reducing retain and test accuracy. Thus, this algorithm produces a certificate through which test-time privacy can be verified while still obtaining a good privacy-utility tradeoff. For both algorithms, tables are included in Appx. J for completeness.

Pareto Frontiers: To better understand the structure of our problem, we explore the Pareto frontier in Fig. 3. We observe that for MNIST, CIFAR10, and CIFAR100, various θ can provide good retain accuracy, albeit at the cost of uniformity. In general, we find that $\theta \approx 0.75$ offers a solid privacy-utility tradeoff. Thus, the γ in Eq. (5) can be chosen fairly large while ensuring low confidence distance.

Additional Experiments: We conduct various additional experiments in in Appx. J and briefly comment about them here. Firstly, we obtain excellent performance for TinyImageNet, ImageNet, and ViT in Appx. J.3; for a ViT pretrained on ImageNet, we obtain a confidence distance of 0.037 with a retain and test accuracy drop of $< 1\%$. Secondly, in Appx. J.4, we show that Alg. 1 obtains high confidence distances on the retain and test sets versus LabelDP; we study additional baselines in Appx. J.12. Thirdly, we study the optimization dynamics of Alg. 1 in Appx. J.6, providing mathematical and empirical evidence for the necessity of early stopping in large models when using Alg. 1. Fourthly, we evaluate our method on several strong TTP attacks, demonstrating that we can still offer effective defense, especially when compared to pretraining or retraining, in Appx. J.7. Fifthly, in Appx. J.8, we find that we preserve strong accuracy and high confidence, as desired, on test instances which are nearest neighbors to the protected set instances. Thus, an adversary querying nearby instances outside of the protected set does not suffice to circumvent our algorithms. Sixthly, we find that we can induce uncertainty on protected instances which were not part of the original training dataset, while still preserving retain and test accuracies, in Appx. J.9. Seventhly, we provide ablations on the size of our protected set in Appx. J.10. Finally, we compare our confidence metric to an ℓ_2 uniformity metric, finding that they highly correlate, in Appx. J.11.

7 Discussion

We present *test-time privacy*, a threat model in which an adversary seeks to directly use a confident prediction for harm. This contrasts with existing work like PATE and LabelDP, which focus on protecting against model inversion and leakage of ground truth labels. To protect against a test-time privacy adversary, we

present multiple algorithms to induce uniformity on a known corrupted subset while preserving utility on the rest of the data instances. This can be used to prevent adversaries from taking advantage of model outputs. Furthermore, we prove a privacy-utility tradeoff for our algorithms, providing a tight bound which is empirically verified. We hope our test-time privacy can further inspire the community to explore different threat models for sensitive data. Limitations and future directions are provided in Appx. E.

References

Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. *Proceedings of the Conference on on Computer and Communications Security (CCS)*, pp. 308–318, 2016.

Naman Agarwal, Brian Bullins, and Elad Hazan. Second-order stochastic optimization in linear time. *stat*, 1050:15, 2016.

Mohammad Al-Rubaie and J Morris Chang. Privacy-preserving machine learning: threats and solutions. *IEEE Security & Privacy*, 17(2):49–58, 2019.

Youssef Allouah, Joshua Kazdan, Rachid Guerraoui, and Sanmi Koyejo. The utility and complexity of in-and out-of-distribution machine unlearning. *International Conference on Machine Learning (ICML)*, 2025.

Kareem Amin, Alex Kulesza, and Sergei Vassilvitskii. Practical considerations for differential privacy. *arXiv preprint arxiv:2408.07614*, 2024.

Anastasios N Angelopoulos, Michael I Jordan, and Ryan J Tibshirani. Gradient equilibrium in online learning: theory and applications. *arXiv preprint arXiv:2501.08330*, 2025.

Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shiho Moriai, et al. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5): 1333–1345, 2017.

Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: analytical calibration and optimal denoising. *International Conference on Machine Learning (ICML)*, 2018.

Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. *IEEE Symposium on Security and Privacy (SP)*, pp. 141–159, 2021.

Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *Transactions on Computation Theory (TOCT)*, 6(3):1–36, 2014.

Mark Bun, Damien Desfontaines, Cynthia Dwork, Naor Moni, Kobbi Nissim, Aaron Roth, Adam Smith, Thomas Steinke, Jonathan Ullman, and Salil Vadhan. Statistical inference is not a privacy violation, 2021.

Robert Istvan Busa-Fekete, Umar Syed, Sergei Vassilvitskii, et al. On the pitfalls of label differential privacy. In *NeurIPS Workshops*, 2021.

Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Eric Wallace, David Rolnick, and Florian Tramèr. Stealing part of a production language model. In *International Conference on Machine Learning (ICML)*, 2024.

Sungmin Cha, Sungjun Cho, Dasol Hwang, Honglak Lee, Taesup Moon, and Moontae Lee. Learning to unlearn: instance-wise unlearning for pre-trained classifiers. *AAAI Conference on Artificial Intelligence*, 38 (10):11186–11194, 2024.

Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. *Advances in Neural Information Processing Systems (NeurIPS)*, 21, 2008.

Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(3), 2011.

Lynn Chua, Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Amer Sinha, and Chiyuan Zhang. Scalable dp-sgd: shuffling vs. poisson subsampling. *Advances in Neural Information Processing Systems (NeurIPS)*, 37:70026–70047, 2024.

Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Gianna M Del Corso, Antonio Gulli, and Francesco Romani. Ranking a stream of news. In *Proceedings of the 14th international conference on World Wide Web*, pp. 97–106, 2005.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing*, 29(6):141–142, 2012.

Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. *Proceedings of the Symposium on Principles of Databases (PODS)*, pp. 202–210, 2003.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*, 2021.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. *Proceedings of the Theory of Cryptography Conference (TCC)*, 2006.

Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3-4):211–407, 2014.

Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the Conference on on Computer and Communications Security (CCS)*, pp. 1322–1333, 2015.

GDPR. Regulation (EU) 2016/679 of the European Parliament and of the Council. OJ L 119, 4.5.2016, p. 1–88, 2016. URL https://data.europa.eu/eli/reg/2016/679/oj.

Badih Ghazi, Noah Golowich, Ravi Kumar, Pasin Manurangsi, and Chiyuan Zhang. Deep learning with label differential privacy. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:27131–27145, 2021.

Shafi Goldwasser, Silvio Micali, and Chales Rackoff. The knowledge complexity of interactive proof-systems. In *Providing sound foundations for cryptography: On the work of shafi goldwasser and silvio micali*, pp. 203–225. ACM Books, 2019.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)*, 2015.

Google. Vision transformer pretrained on imagenet-21k, 2023. URL https://huggingface.co/google/vit-base-patch16-224.

Laura Graves, Vineel Nagisetty, and Vijay Ganesh. Amnesiac machine learning. In *AAAI Conference on Artificial Intelligence*, 2021.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. *International Conference on Machine Learning (ICML)*, 2017.

Niv Haim, Gal Vardi, Gilad Yehudai, michal Irani, and Ohad Shamir. Reconstructing training data from trained neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.

Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations (ICLR)*, 2019.

C-L Hwang and Abu Syed Md Masud. *Multiple objective decision making—methods and applications: a state-of-the-art survey*, volume 164. Springer Science & Business Media, 2012.

International Association of Privacy Professionals. Swedish court rejects google's appeal in RTBF case, 2020. URL https://iapp.org/news/a/swedish-court-rejects-googles-appeal-in-rtbf-case.

Gautam Kamath. Lecture 12: what is privacy? *Lectures on private ml and stats*, 2020. URL http://www.gautamkamath.com/CS860notes/lec12.pdf.

Anastasia Koloskova, Youssef Allouah, Animesh Jha, Rachid Guerraoui, and Sanmi Koyejo. Certified unlearning for neural networks. *International Conference on Machine Learning (ICML)*, 2025.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical Report from the University of Toronto*, 2009.

Solomon Kullback and Richard A Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded machine unlearning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Ya Le and Xuan S. Yang. Tiny imagenet visual recognition challenge, 2015.

Joon-Woo Lee, HyungChul Kang, Yongwoo Lee, Woosuk Choi, Jieun Eom, Maxim Deryabin, Eunsang Lee, Junghyun Lee, Donghoon Yoo, Young-Sik Kim, et al. Privacy-preserving machine learning with fully homomorphic encryption for deep neural network. *IEEE Access*, 10:30039–30054, 2022.

Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. In *International Conference on Learning Representations (ICLR)*, 2018.

Ninghui Li, Wahbeh Qardaji, and Dong Su. On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. *Proceedings of the Conference on on Computer and Communications Security (CCS)*, pp. 32–33, 2012.

Hui Liu, Yibo Dou, Kai Wang, Yunmin Zou, Gan Sen, Xiangtao Liu, and Huling Li. A skin disease classification model based on multi scale combined efficient channel attention module. *Scientific Reports*, 15(1):6116, 2025.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations (ICLR)*, 2018.

Günter Mayer. On the convergence of the neumann series in interval analysis. *Linear algebra and its applications*, 65:63–70, 1985.

Frank Mcsherry. Statistical inference considered harmful, 2016. URL https://github.com/frankmcsherry/blog/blob/master/posts/2016-06-14.md.

Ronak Mehta, Sourav Pal, Vikas Singh, and Sathya N Ravi. Deep unlearning via randomized conditionally independent hessians. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

Merriam-Webster. Privacy. *Merriam-Webster Dictionary*, 2022. URL https://www.merriam-webster.com/dictionary/privacy.

Microsoft. Resnet50 pretrained on imagenet-21k, 2024. URL https://huggingface.co/microsoft/resnet-50.

Kaisa Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media, 1999.

Siqiao Mu and Diego Klabjan. Rewind-to-delete: certified machine unlearning for nonconvex functions. *arXiv preprint arXiv:2409.09778*, 2025.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. *NeurIPS Workshops*, pp. 4, 2011.

Jorge Nocedal and Stephen J Wright. *Numerical optimization.* Springer, 1999.

Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. Scalable private learning with pate. *International Conference on Learning Representations (ICLR)*, 2018.

Panos M Pardalos, Antanas Žilinskas, Julius Žilinskas, et al. *Non-convex multi-objective optimization.* Springer, 2017.

Tim Pearce, Alexandra Brintrup, and Jun Zhu. Understanding softmax confidence and uncertainty. *arXiv preprint arXiv:2106.04972*, 2021.

Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.

Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *International Conference on Learning Representations (ICLR)*, 2017.

Mark S Pinsker. Information and information stability of random variables and processes. *Holden-Day*, 1964.

Xinbao Qiao, Meng Zhang, Ming Tang, and Ermin Wei. Hessian-free online certified unlearning. *International Conference on Learning Representations (ICLR)*, 2025.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015.

Ahmed Salem, Giovanni Cherubin, David Evans, Boris Köpf, Andrew Paverd, Anshuman Suri, Shruti Tople, and Santiago Zanella-Béguelin. Sok: Let the privacy games begin! a unified treatment of data inference privacy in machine learning. In *2023 IEEE Symposium on Security and Privacy (SP)*, pp. 327–345. IEEE, 2023.

Stefan Schoepf, Michael Curtis Mozer, Nicole Elyse Mitchell, Alexandra Brintrup, Georgios Kaissis, Peter Kairouz, and Eleni Triantafillou. Redirection for erasing memory (rem): Towards a universal unlearning method for corrupted data. *arXiv preprint arXiv:2505.17730*, 2025.

Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: algorithms for machine unlearning. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:18075–18086, 2021.

Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. *IEEE Symposium on Security and Privacy (SP)*, pp. 3–18, 2017.

Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. Machine learning models that remember too much. *Proceedings of the Conference on on Computer and Communications Security (CCS)*, pp. 587–601, 2017.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations (ICLR)*, 2021.

Jingwei Sun, Ang Li, Binghui Wang, Huanrui Yang, Hai Li, and Yiran Chen. Soteria: Provable defense against privacy leakage in federated learning from representation perspective. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

Xiaoxiao Sun, Jufeng Yang, Ming Sun, and Kai Wang. A benchmark for automatic visual classification of clinical skin disease images. *European Conference on Computer Vision (ECCV)*, 2016.

TorchVision. Torchvision: Pytorch's computer vision library, 2016. URL https://github.com/pytorch/vision.

Cuong Tran and Ferdinando Fioretto. Personalized privacy auditing and optimization at test time. *arXiv preprint arXiv:2302.00077*, 2023.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.

Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

Zhibo Wang, Mengkai Song, Zhifei Zhang, Yang Song, Qian Wang, and Hairong Qi. Beyond inferring class representatives: User-level privacy leakage from federated learning. *IEEE Conference on Computer Communications*, 2019.

Annie White. DMVs can (and do) collect and sell your personal data. *Car and Driver*, 2020. URL https://www.caranddriver.com/features/a32035408/dmv-selling-driver-data/.

Ruihan Wu, Jin Peng Zhou, Kilian Q Weinberger, and Chuan Guo. Does label differential privacy prevent label inference attacks? *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023.

Taihong Xiao, Yi-Hsuan Tsai, Kihyuk Sohn, Manmohan Chandraker, and Ming-Hsuan Yang. Adversarial learning of privacy-preserving and task-oriented representations. *AAAI Conference on Artificial Intelligence*, 2020.

Jinghan Yang and Jiayu Weng. Detecting mislabeled and corrupted data via pointwise mutual information. *arXiv preprint arXiv:2508.07713*, 2025.

Jufeng Yang, Xiaoxiao Sun, Jie Liang, and Paul L Rosin. Clinical skin lesion diagnosis using representations inspired by dermatologist criteria. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. *2018 IEEE Symposium on Computer Security Foundations (CSF)*, pp. 268–282, 2018.

Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. Differentially private fine-tuning of language models. *International Conference on Learning Representations (ICLR)*, 2022.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 5824–5836, 2020.

Binchi Zhang, Yushun Dong, Tianhao Wang, and Jundong Li. Towards certified unlearning for deep neural networks. *International Conference on Machine Learning (ICML)*, 2024.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations (ICLR)*, 2017.

Rui Zhang, Song Guo, Junxiao Wang, Xin Xie, and Dacheng Tao. A survey on gradient inversion: Attacks, defenses and future directions. *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2022.

Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. idlg: Improved deep leakage from gradients. *arXiv preprint arXiv:2001.02610*, 2020.

Kairan Zhao, Meghdad Kurmanji, George-Octavian Bărbulescu, Eleni Triantafillou, and Peter Triantafillou. What makes unlearning hard and what to do about it. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Mingyi Zhou, Xiang Gao, Jing Wu, John Grundy, Xiao Chen, Chunyang Chen, and Li Li. Modelobfuscator: obfuscating model information to protect deployed ml-based systems. *Proceedings of the Symposium on Software Testing and Analysis*, pp. 1005–1017, 2023.

Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Appendix

Table of Contents

A Defining Test-Time Privacy Attacks

In what follows, in light of our threat model provided in Sec. 3, we design some simple test-time privacy attacks to motivate our problem. We also include experiments on these attacks, and how Alg. 1 performs against them, in Appx. J.7.

Our first simple algorithm is to add a small amount of uniformly sampled Gaussian noise, presented in Alg. 4. We find that this is not very effective in increasing confidence distance, as demonstrated in Tab. J.8 and Tab. J.9. When it brings the confidence distances from low to moderate, the model is usually confidently wrong, as demonstrated in Tab. J.10.

One way to more optimally attack the TTP of a pretrained model is by finding instances in a ϕ-ball around the protected set instances that maximize the prediction confidence. To design such an attack, suppose we have a pretrained classifier $f_{\boldsymbol{w}^*} : \mathcal{X} \rightarrow \Delta_{|\mathcal{Y}|}$. Here, $f_{\boldsymbol{w}}(\boldsymbol{x}) = \text{softmax}(\boldsymbol{z}(\boldsymbol{x}))$, for $\boldsymbol{x} \in \mathcal{X}$, where $\boldsymbol{z}$ is a vector of logits. For a protected set instance, we begin by adding a small amount of uniform noise to break symmetry and obtain a nonzero gradient. We then want to obtain the worst-case perturbation over the logits by solving the optimization problem:

$$\max_{\phi} \max_{j} \boldsymbol{z}_{\boldsymbol{w}^*}^{(j)}(\boldsymbol{x} + \phi), \tag{A.7}$$

$$\text{s.t. } ||\phi||_{\infty} \leq \gamma. \tag{A.8}$$

Since the max function is not differentiable everywhere, we use LogSumExp to approximate it. Denote $\rho(f_{\boldsymbol{w}}(\boldsymbol{x} + \boldsymbol{\phi})) = \log \sum_{j=1}^{|\mathcal{Y}|} \exp\left(\boldsymbol{z}_{\boldsymbol{w}}^{j}(\boldsymbol{x} + \boldsymbol{\phi})\right)$. This yields the optimization problem:

$$\max_{\phi} \rho(f_{\boldsymbol{w}^*}(\boldsymbol{x} + \delta)), \tag{A.9}$$

$$\text{s.t. } ||\phi||_{\infty} \leq \gamma. \tag{A.10}$$

Following the Fast Gradient Sign Method (FGSM), a simple attack used to generate adversarial examples (Goodfellow et al., 2015), we design an attack as Alg. 5. Intuitively, we take a single linear step towards maximizing the function. We design also design stronger attack based on Projected Gradient Descent (PGD) (Madry et al., 2018) as Alg. 6, taking 40 steps while incrementally maximizing the confidence function while projecting back to the ball around the original instance. Empirical results are in Appx. J.7.

B Additional Related Work

Differential Privacy: Differential privacy has widely been studied in the ML community in order to ensure privacy-preservation (Chaudhuri & Monteleoni, 2008); (Chaudhuri et al., 2011); (Abadi et al., 2016); (Chua et al., 2024). There also exist methods to finetune pretrained models to satisfy differential privacy (Yu et al., 2022). Furthermore, there are also ways to aggregate label noise to preserve privacy (Papernot et al., 2018).

However, differential privacy is designed to address an entirely different threat model than ours. In particular, in the threat model of differential privacy, an adversary seeks to use model outputs to recover private information about data instance $\boldsymbol{x}_p$ corresponding to person p with e.g. a model inversion attack. A differentially private classifier generally results in confident, accurate predictions. This does not address our threat model, where an adversary may use confident model outputs to violate the privacy of person p in a different manner, taking advantage of them directly to use as a warrant to cause harm to person p.

Label Differential Privacy: Similarly, our formulation differs from label differential privacy (LabelDP) (Ghazi et al., 2021), which seeks to protect an adversary from learning the true labels of the instances in the training data. Given an instance, even after computing $f(\boldsymbol{x}_p)$, under LabelDP an adversary cannot be confident that $f(\boldsymbol{x}_p) = y$. However, LabelDP is applied to the entire dataset; our threat model involves only

a particular subset of the training data. Furthermore, we do not need to protect the user's ground truth label, necessarily. In our law enforcement example in Sec. 1, the agency does not care about the ground truth label. Instead, they want any confirmation such that they have a warrant to act adversarially towards person p; for this, a confident prediction by model f suffices. Finally, LabelDP results in poor retain and test accuracy for larger datasets e.g. CIFAR100, as demonstrated in Fig. 2.

Furthermore, from the perspective of protecting the privacy of the labels themselves, rather than protecting against any confident prediction, Busa-Fekete et al. (2021) demonstrate that testing a model, trained with LabelDP, on the training dataset allows an adversary to recover the labels of the label-private data with high probability. Since our algorithms induce uniformity, an adversary cannot infer the correct protected set labels by testing the model on the training dataset; thus, we provide better privacy against this threat model than LabelDP as well. Wu et al. (2023) argue that, under this threat model where one seeks to protect the labels, any model that generalizes must leak the accurate labels when tested on the training data. However, as we demonstrate by inducing uniformity while maintaining high test accuracy, this only holds when the model is to be tested on the *entire* training data, not a *subset* of the training data (or other test instances which are known to be corrupted), as in our setting.

Label Model Inversion Attacks: Related to LabelDP are model inversion attacks to recover the ground truth labels, like gradient inversion (Zhang et al., 2022); (Zhu et al., 2019); (Zhao et al., 2020). Yet, these methods do not report the confidence values for the recovered labels. Thus, they do not constitute test-time privacy attacks within our threat model. Furthermore, by the same token as above, an adversary seeks to recover a confident prediction to use as a warrant, not necessarily the ground truth labels. Still, these methods could potentially be extended to test-time privacy attacks by reporting a confidence score for the recovered labels. We leave this to future work.

Other Paradigms in Privacy: Other paradigms in the privacy literature correspond to a notion of "test-time privacy" which differ from our threat model. For example, several works study defense against model inversion attacks as test-time privacy (Wang et al., 2019); (Xiao et al., 2020); (Sun et al., 2021); (Tran & Fioretto, 2023). However, this is a separate threat model from ours; the adversary already has access to the instance $\boldsymbol{x}_p$ within our threat model.

Misclassification & Relabeling in Machine Unlearning: Recently, methods have emerged to finetune a model to misclassify rather than mimicking retraining from scratch (Cha et al., 2024). There are other similar relabeling methods in the debiasing literature which could be used for this purpose (Angelopoulos et al., 2025). However, these methods often achieve poor performance on the remaining training data and fail to provide protection against our threat model in all cases, as described in Sec. 3. For example, in the insurance example in Sec. 1, it may still be problematic to classify the user as "Benign" instead of "Melanoma"; for example, the user of model f (i.e., the TTP adversary ν) could be a medical professional instead of an insurance provider.

Model Calibration and Confidence: In our setting, we use the model softmax outputs to represent the adversary's confidence in the final prediction. However, some argue that this type of interpretation is incorrect, i.e. ML models are poorly calibrated (Guo et al., 2017). Still, this interpretation is common (Pearce et al., 2021), and thus a model user would likely rely on the softmax outputs as the confidence scores. We leave to inducing uncertainty over the calibrated outputs to future work.

Out-of-Distribution and Data Abstention: Notably, our threat model cannot be addressed by approaches based on detecting or flagging corrupted data via out-of-distribution and data abstention techniques. This is because we assume a worst-case adversary that may, due to negligence or maliciousness, seek to obtain a harmful prediction regardless if corruption is flagged to them or not. Furthermore, the corrupted instance is not necessarily from a different distribution than the clean instance. In particular, out-of-distribution (OOD) methods (Hendrycks et al., 2019; Lee et al., 2018) assume this different setting.

C Uniformity Metric

The confidence distance quantifies the adversary's confidence in their final prediction, i.e. the difference between the argmax softmax score and the uniform softmax score. Importantly, our method aims to have the

adversary lack confidence in their final prediction. Thus, our metric captures what we aim to measure and is interpretable, since it is minimized at 0.

Furthermore, confidence distance allows us to quantify how uncertain the model is without relying on accuracy, since a drop in protected set accuracy is not the goal of our formulation. Next, if the maximum confidence score is very close to the uniform distribution, the probability mass of the output distribution must be distributed over the other softmax outputs, clearly yielding that the higher our uniformity metric, the more confident our model is, and the lower our uniformity metric, the less confident our model is. Additionally, it takes the dataset into account; for example, in CIFAR10, one would expect a uniformity score of ≈ 0.2 to be reasonable, as then the adversary can only be $\approx 30\%$ confident that they have a useful prediction. However, for CIFAR100, a uniformity score of ≈ 0.2 is much better, as it implies that an adversary can only be $\approx 21\%$ confident that they have a useful prediction.

One objection to the use of this metric may be that it does not indicate uniformity if it is low. For example, on CIFAR10, one could have a confidence score of 0.25, which yields that the max softmax output is 0.35. There could be three other nonzero softmax outputs of 0.3, 0.2, 0.1, 0.05; this clearly is not uniform. However, this ensures test-time privacy; a test-time privacy adversary now has little confidence in their prediction, even if they choose the first one, rendering their warrant for misuse of sensitive data useless.

We empirically compare our confidence distance metric to other similar metrics in Appx. J, finding that when our confidence metric is minimized, other metrics are minimized.

D   Test-Time Privacy Examples

Here, we provide a set of examples of the TTP threat model:

Health Insurance: Suppose an open-weight medical imaging model f is released, designed to perform multiclass classification of skin photos into categories like "Dysplastic Nevus", "Benign Keratosis", which are usually harmless, or serious classes like "Melanoma" (Sun et al., 2016). A person p posts a photo of a harmless birthmark on his arm to a public health forum to ask a question. During the upload, an e.g. server error or compression issue causes the image file to become corrupted, severely distorting the birthmark. This results in a photo $\boldsymbol{x}_p$. Next, a health insurance startup decides to build risk profiles by scraping these public forums. They download the open-weight model f to automatically screen images for potential health liabilities. When they feed $\boldsymbol{x}_p$ into f, it confidently classifies $\boldsymbol{x}_p$ as "Melanoma". This erroneous classification is then automatically added to person p's risk profile, resulting person p being unfairly denied coverage.

Criminal Records: Suppose a model f is trained on criminal records to predict individual crime likelihood. Additionally, suppose the criminal record $\boldsymbol{x}_p$ of a person p is corrupted and publicly available. Then, $f(\boldsymbol{x}_p)$ predicts that person p is highly likely to commit crime. An adversarial law enforcement agency, or even a prospective employer, may ignore or be unaware of warnings about the data being corrupted, rendering a dangerous scenario for person p.[4] To make this clear, provide a figure similar to that of Fig. 1 at Fig. D.4.

Mortgage Loans: Suppose a model f is trained on various items relevant to whether one receives a mortgage loan or not, like bank statements and past rent payments. Person p has corrupted rent payment history $\boldsymbol{x}_p$. Then, the bank runs model f and obtains $f(\boldsymbol{x}_p)$, which confidently says that $\boldsymbol{x}_p$ is undeserving of a loan.

Car Insurance: Suppose a model f is trained on one's history of car accidents. Person p has corrupted car accident history $\boldsymbol{x}_p$. Then, when applying for car insurance, the provider runs model f and obtains $f(\boldsymbol{x}_p)$, which confidently says that $\boldsymbol{x}_p$ is undeserving of a loan.[5]

We provide an additional example in the generative setting as well:

News Articles: Consider a text-to-image generative model trained on a large dataset, including web data, which has web articles and associated images. A popular news site publishes an article about a businessperson,

---

[4]Recently, ML model providers have been involved in privacy cases involving criminal records (International Association of Privacy Professionals, 2020), making this threat pertinent.

[5]Note that recent, the Department of Motor Vehicles in America has been selling driving records, making this threat pertinent (White, 2020).

Figure D.4: An adversary, like an employer (🏢), can query a pretrained model f (⬤) and use its outputs to make harmful decisions. However, after running our algorithm, the new model f_U (⬤) provides maximal uncertainty, protecting against such an adversary. This is a duplicate of Fig. 1, to make clear how TTP extends to other settings.

but mistakenly uses a picture of an unrelated individual p, $\boldsymbol{x}_p$, as the header image. This creates a strong, albeit false, association between this person's likeness and the (perhaps negative) content of the article. When prompted with a string similar to the headline of the news article, the model generates an image (or a similar image) of person p, algorithmically cementing a false narrative about person p.

E   Limitations and Future Directions

Notably, our presented method only applies to classification. While we provide an example of TTP in the generative setting in Appx. D, we leave extending TTP to e.g. diffusion models (Song et al., 2021) to future work. This extension is nontrivial, since there are uncountably infinitely many instances to induce uniformity over. Furthermore, following the newspaper example given in Appx. D, the "protected set" is often a *concept* or *association*, not a discrete set of instances. This holds similarly for sequence-to-sequence generation (Vaswani et al., 2017). Furthermore, even in the discriminative setting, we focus our method on image classification. Extending our methods to the text setting, which is nontrivial due to discrete inputs, remains as future work.

From an algorithmic perspective, in Alg. 1, we use linear scalarization to design our objective (Hwang & Masud, 2012). One can instead design an objective using ε-constraints (Miettinen, 1999), which can then be solved by an augmented Lagrangian method (Nocedal & Wright, 1999). Furthermore, our certified algorithms, especially Alg. 3, are very sensitive to hyperparameters. Specifically, λ and J must be nearly precisely $|\lambda_{\min}|$ and $|\lambda_{\max} + \lambda|$. Otherwise, the estimator diverges aggressively. Furthermore, other Hessian estimators (Mehta et al., 2022) do not admit bounds which are necessary for our certified Gaussian mechanism. Finally, taking λ large to ensure that the Hessian is p.d. and hence invertible, i.e. local convex approximation, is problematic in our multiobjective setting since it results in a step off the Pareto frontier. These limitations can be overcome by adopting recent Hessian-free approaches from certified unlearning (Qiao et al., 2025; Mu & Klabjan, 2025; Koloskova et al., 2025). However, adopting these approaches is nontrivial, since they rely on assumptions about the retrain-from-scratch (target) model which do not hold for the Pareto model, e.g. the structure of its gradient updates. As such, we leave this to future work. Additionally, we only use a small subset of the training dataset as a protected set, in line with our assumption that falsified public records are rare. While we provide an ablation on this in Appx. J.10, we leave algorithms which can systematically induce uncertainty on large protected sets while still keeping strong retain accuracy to future work.

Next, we do not study the case where the model provider is potentially adversarial as well. Instead, we assume that both the model provider and person p are benign. In the setting where the model provider is also adversarial, they may induce uncertainty on incriminating instances. In that same vein, person p may submit an uncorrupted instance to the model provider, causing the model provider to mistakenly induce uncertainty. While verifying the protected set is out of scope for our work, zero-knowledge (Goldwasser et al., 2019) can be useful for a trusted third-party to verify the construction of the protected set.

Additionally, while in our threat model person p's identifying information is publicly attached to $\boldsymbol{x}_p$, inducing uncertainty does increase the success of membership inference attacks. In the case that the adversary releases a new model, this could constitute a threat to person p. We leave this separate threat model to future work. In that same vein, we leave a guarantee on the certified Pareto learner (Def. 4.4) that is fully tamper-resistant across adversarial actors to future work as well.

Finally, while we proved an upper bound on the retain accuracy drop for our Pareto learner without assuming convexity in Sec. 5, the general theoretical qualities of TTP still remain to be studied. In particular, one could prove lower bounds on the generalization error or the sample complexity of TTP. This would allow us to design more principled TTP algorithms which e.g. achieve the minimax rate. However, obtaining such lower bounds is nontrivial. For example, current approaches which prove lower bounds for deletion capacity (sample complexity) in the related setting of certified unlearning (Allouah et al., 2025) rely heavily on the fact that (ε, δ)-unlearning aims to mimic $\mathcal{A}(\mathcal{D}_r)$, whereas (ε, δ)-TTP aims to mimic $\mathcal{M}_\theta(\mathcal{D})$. Furthermore, these works assume strongly convex losses, making them inapplicable to the deep learning setting. We thus leave this to future work.

F   Designing Certified Algorithms

In what follows, we design $(\varepsilon, \delta, \theta)$-certified Pareto learners. A symbol table can be found at Appx. K.

In our setting, the original model is obtained using ERM over some loss function $\mathcal{L}_\mathcal{A}$, some dataset $\mathcal{D}$, and some parameter space $\mathcal{W}$. Furthermore, we consider the common scenario where the cumulative loss $\mathcal{L}_\mathcal{A}$ over the dataset is a finite sum of individual losses $\ell_\mathcal{A}$. Thus, we denote the pretrained model as:

$$\boldsymbol{w}^* = \mathcal{A}(D) := \arg\min_{\boldsymbol{w} \in \mathcal{W}} \mathcal{L}_\mathcal{A}(\boldsymbol{w}, \mathcal{D}) = \arg\min_{\boldsymbol{w} \in \mathcal{W}} \sum_{i=1}^{|\mathcal{D}|} \ell_\mathcal{A}(\boldsymbol{w}, \mathcal{D}^{(i)}). \tag{F.11}$$

By Prop. H.10, we can similarly obtain a uniform learner through ERM with respect to some loss function $\mathcal{L}_\mathcal{K}$. Furthermore, in our setting, we have the protected set $\mathcal{D}_p$ and retain set $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_p$. Thus, the uniform learner over the protected set can be characterized as:

$$\mathcal{K}(\mathcal{D}_p) := \arg\min_{\boldsymbol{w} \in \mathcal{W}} \mathcal{L}_\mathcal{K}(\boldsymbol{w}, \mathcal{D}) = \sum_{i=1}^{|\mathcal{D}_p|} \ell_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p^{(i)}). \tag{F.12}$$

Let $\theta \in (0, 1)$ be a tradeoff parameter between uniformity over the protected set and utility over the retain set. This yields a concrete characterization of $\mathcal{M}_\theta$ as:

$$\tilde{\boldsymbol{w}}^* = \mathcal{M}_\theta(D) := \arg\min_{\boldsymbol{w} \in \mathcal{W}} \theta \mathcal{L}_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p) + (1 - \theta)\mathcal{L}_\mathcal{A}(\boldsymbol{w}, \mathcal{D}_r), \tag{F.13}$$

$$= \arg\min_{\boldsymbol{w} \in \mathcal{W}} \theta \sum_{i=1}^{|\mathcal{D}_p|} \ell_\mathcal{K}(\boldsymbol{w}, \mathcal{D}_p^{(i)}) + (1 - \theta) \sum_{i=1}^{|\mathcal{D}_r|} \ell_\mathcal{A}(\boldsymbol{w}, \mathcal{D}_r^{(i)}). \tag{F.14}$$

To design an algorithm which takes in $\mathcal{D}, \mathcal{D}_r$, and $\boldsymbol{w}^*$ and outputs a parameter which satisfies Def. 4.4, we follow the methodology of certified unlearning Zhang et al. (2024), which seeks to satisfy Def. 4.3.

First, we simplify the problem of deriving a model that satisfies Def. 4.4:

Theorem F.1. *(Certification Guarantee) Let $\tilde{\boldsymbol{w}} := \mathcal{F}(\mathcal{D}, \mathcal{D}_p, \boldsymbol{w}^*)$ be an approximation to $\tilde{\boldsymbol{w}}^*$. Suppose $||\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^*||_2 \leq \Delta$. Then, $\mathcal{U}(\mathcal{D}, \mathcal{D}_p, A(D)) = \boldsymbol{w}^- = \tilde{\boldsymbol{w}} + \boldsymbol{Y}$ is a $(\epsilon, \delta, \theta)$ certified uniformity algorithm, where $\boldsymbol{Y} \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I})$ and $\sigma \geq \frac{\Delta}{\epsilon}\sqrt{2\ln(1.25/\delta)}$.*

Proof: See Appx. H.4.

Thus, it then suffices to find an approximation of $\tilde{\boldsymbol{w}}^*$, i.e. a form for $\mathcal{F}(\mathcal{D}, \mathcal{D}_p, \boldsymbol{w}^*)$ and its associated Δ. To do so, we consider the two assumptions Asm. 5.3 and Asm. 5.4.

For any $\boldsymbol{w} \in \mathcal{W}$, denote $\nabla_{\boldsymbol{w}, \mathcal{K}, \mathcal{A}} := \nabla_{\boldsymbol{w}}(\theta\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r))$, the gradient of the objective of $\mathcal{M}_\theta$ with respect to $\boldsymbol{w}$, and $\boldsymbol{H}_{\boldsymbol{w}, \mathcal{K}, \mathcal{A}} := \nabla_w^2(\theta\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r))$, the Hessian of the objective of $\mathcal{M}_\theta$ with respect to $\boldsymbol{w}$. We thus have $\nabla_{\boldsymbol{w}, \mathcal{K}, \mathcal{A}} = \theta\nabla_{\boldsymbol{w}, \mathcal{K}} + (1-\theta)\nabla_{\boldsymbol{w}, \mathcal{A}}$, and similarly for the Hessian.

Next, letting $g(\boldsymbol{w}) := \nabla_{\boldsymbol{w}, \mathcal{K}, \mathcal{A}}$, by Taylor's theorem, expanding $g(\tilde{\boldsymbol{w}}^*)$ around $\boldsymbol{w}^*$, we have that:

$$g(\tilde{\boldsymbol{w}}^*) \approx g(\boldsymbol{w}^*) + Dg|_{\boldsymbol{w}^*}(\tilde{\boldsymbol{w}}^* - \boldsymbol{w}^*). \tag{F.15}$$

.

Note that $g(\tilde{\boldsymbol{w}}^*) = 0$, since $\tilde{\boldsymbol{w}}^*$ is the minimizer of the objective in $\mathcal{M}_\theta$. Isolating $\tilde{\boldsymbol{w}}^*$ and using the definition of g, we then have that:

$$\tilde{\boldsymbol{w}}^* \approx \boldsymbol{w}^* - \boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}^{-1}\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}. \tag{F.16}$$

.

Thus, we let $\tilde{\boldsymbol{w}} = \boldsymbol{w}^* - \boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}^{-1}\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$. This yields the following general form of Δ:

Proposition F.2. *Suppose Asm. 5.3 and Asm. 5.4 hold. Suppose $\tilde{\boldsymbol{w}} = \boldsymbol{w}^* - \boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}^{-1}\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$. Then,*

$$||\tilde{\boldsymbol{w}}^* - \tilde{\boldsymbol{w}}||_2 \leq \frac{\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}}}{2}||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}^{-1}||_2||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2^2. \tag{F.17}$$

Proof: See Appx. H.5.

We then use local convex approximation (Nocedal & Wright, 1999) to bound $||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}^{-1}||_2$. To that end, we let the objective of $\mathcal{M}_\theta$ have a regularization term $\frac{\lambda}{2}||\boldsymbol{w}||_2^2$, yielding the inverse Hessian $||(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda\boldsymbol{I})^{-1}||_2$; thus, in Prop. F.2, the norm of the inverse Hessian is replaced by $||(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda\boldsymbol{I})^{-1}||_2$. It then suffices to bound this term.

Additionally, note that since the objective of $\mathcal{M}_\theta$ is nonconvex, the Hessian may not be invertible, i.e. $\lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}) < 0$. However, $\lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda\boldsymbol{I}) = \lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}) + \lambda$. Thus, for λ sufficiently large, we can make $\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda\boldsymbol{I}$ positive definite and hence invertible, resolving this issue. In particular, we can take $\lambda > ||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2$.

Furthermore, we let $||\boldsymbol{w}||_2 \leq C$ in $\mathcal{M}_\theta$ and $\mathcal{A}$, i.e. $\mathcal{M}_\theta = \arg\min_{||\boldsymbol{w}||_2 \leq C, \boldsymbol{w} \in \mathcal{W}} \theta\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_{\mathcal{A}}) + \frac{\lambda}{2}||\boldsymbol{w}||_2^2$ and $\mathcal{A}(\mathcal{D}) = \arg\min_{\boldsymbol{w} \in \mathcal{W}, ||\boldsymbol{w}|| \leq C} \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D})$. Note that, as mentioned in (Zhang et al., 2024), unlearning methods implicitly assume this.

Together, these two methods yield a tractable form of Δ:

Proposition F.3. *Suppose Asm. 5.3 and Asm. 5.4 hold. Suppose $\tilde{\boldsymbol{w}} = \boldsymbol{w}^* - (\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda\boldsymbol{I})^{-1}\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$ where $||\boldsymbol{w}^*||_2, ||\tilde{\boldsymbol{w}}^*||_2 \leq C$. Let $\lambda_{\min} := \lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}})$. Suppose $\lambda > ||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2$. Then,*

$$||\tilde{\boldsymbol{w}}^* - \tilde{\boldsymbol{w}}||_2 \leq \frac{2C((\theta F_K + (1-\theta)F_A)C + \lambda)}{\lambda + \lambda_{\min}}. \tag{F.18}$$

Proof: See Appx. H.6.

While Prop. F.3 does yield a form of $\mathcal{F}$ and Δ, the computation of $\tilde{\boldsymbol{w}}$ requires obtaining the exact inverse Hessian, which has a prohibitively high runtime.

If we consider the additional assumption of convexity, we can take λ very small to ensure the Hessian is invertible, since we have $\lambda_{\min} = 0$. Thus, for convex models e.g. logistic regression with a mean-square uniform loss, this is tractable. This yields Alg. 2.

However, for nonconvex models e.g. large scale neural networks, this is computationally intractable. Thus, to provide better runtime, we derive an asymptotically unbiased estimator of the inverse Hessian. However, the estimator in Zhang et al. (2024) does not trivially extend to our case. In particular, we cannot glean Hessian samples using sampled i.i.d. data from the retain set, because the Hessian in our setting is defined over the protected set as well. Thus, we must derive an unbiased estimator while sampling Hessians from *both* the retain and protected set. As such, following the techniques of (Agarwal et al., 2016), we design an unbiased estimator as follows:

Theorem F.4. *Suppose we have n i.i.d. data samples $(X_1, ..., X_n)$ drawn from $\mathcal{D}_p$ and $\mathcal{D}_r$, uniformly at random, with probabilities θ and $1 - \theta$ respectively. Then, suppose $||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda I||_2 \leq J$. For $t = 1, ..., n$, if $X_t \sim \mathcal{D}_p$ let $\boldsymbol{H}_{t, \lambda} = \boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, t} + \frac{\lambda I}{2\theta}$ and if $X_t \sim \mathcal{D}_r$ let $\boldsymbol{H}_{t, \lambda} = \boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{A}, t} + \frac{\lambda I}{2(1-\theta)}$. Suppose $\lambda > ||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2$. Then, compute:*

$$\tilde{\boldsymbol{H}}_{t,\lambda}^{-1} = \boldsymbol{I} + (\boldsymbol{I} - \frac{\boldsymbol{H}_{t,\lambda}}{J})\tilde{\boldsymbol{H}}_{t-1,\lambda}^{-1}, \ \tilde{\boldsymbol{H}}_{0,\lambda} = \boldsymbol{I}. \tag{F.19}$$

Then, $\frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{J}$ is an asymptotically unbiased estimator for $(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda \boldsymbol{I})^{-1}$

Proof: See Appx. H.7

One simple choice of J is $J = 2\lambda$, by Lemma H.1. However, we let J be free. The computation of the estimator in Theorem F.4, together with Hessian vector product (HVP) techniques (Pearlmutter, 1994), reduces the runtime complexity of our exact algorithm greatly. Additionally, following Agarwal et al. (2016), we can average b unbiased estimators $\frac{\tilde{\boldsymbol{H}}_{t,\lambda}^{-1}}{J}$ as $\frac{1}{b} \sum_{i=1}^{b} \frac{\tilde{\boldsymbol{H}}_{t,\lambda}^{-1,(i)}}{J}$ to achieve better concentration.

Furthermore, we relax the assumption that $\boldsymbol{w}^*$ and $\tilde{\boldsymbol{w}}^*$ are the global minimizers of $\mathcal{L}_{\mathcal{A}}$ and $\theta \mathcal{L}_{\mathcal{A}} + (1 - \theta)\mathcal{L}_{\mathcal{K}}$. We do so because, in practice, it is possible that the data controller trained their model with early stopping, i.e. they did not reach the global minimizer. Altogether, this yields a final form of Δ as:

Theorem F.5. *Let $\tilde{\boldsymbol{w}}^*$ and $\boldsymbol{w}^*$ not be empirical risk minimizers of their respective losses, but rather approximations thereof. Suppose Asm. 5.3 and Asm. 5.4 hold. Suppose $||\boldsymbol{w}^*||_2, ||\tilde{\boldsymbol{w}}^*||_2 \leq C$. Let $\lambda_{\min} := \lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}})$. Suppose $\lambda > ||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2$. Let $\tilde{\boldsymbol{w}} = \boldsymbol{w}^* - \frac{\tilde{\boldsymbol{H}}_{t,\lambda}^{-1}}{J}\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$. Let b be the number of inverse Hessian estimators we average. Letting n be the number of steps taken during unbiased estimation of the inverse Hessian, require $n \geq 2\frac{B}{\lambda + \lambda_{\min}} \ln(\frac{B}{\lambda + \lambda_{\min}} b)$ where $B = \max\{\frac{\theta P_{\mathcal{K}} + \lambda}{|\mathcal{D}_p|}, \frac{(1-\theta)P_{\mathcal{A}} + \lambda}{|\mathcal{D}_r|}\}$. Suppose $||\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2, ||\nabla_{\tilde{\boldsymbol{w}}^*, \mathcal{K}, \mathcal{A}}||_2 \leq G$, With probability larger than $1 - \rho$, we have that:*

$$||\tilde{\boldsymbol{w}}^* - \tilde{\boldsymbol{w}}||_2 \leq \frac{2C((\theta F_{\mathcal{K}} + (1 - \theta)F_{\mathcal{A}})C + \lambda) + G}{\lambda + \lambda_{\min}} \tag{F.20}$$

$$+ (16\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln(\frac{d}{\rho})}{b}} + \frac{1}{16})(2C(\theta P_{\mathcal{K}} + (1 - \theta)P_{\mathcal{A}}) + G). \tag{F.21}$$

where $\zeta_{\min} \geq \min_i \lambda_{\min}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K}, \mathcal{A}}(\boldsymbol{w}, \mathcal{D}^{(i)}))$.

Proof: See Appx. H.8.

Note that if we let $\boldsymbol{w}^*$ be an ERM in Theorem F.5, we can use $\nabla_{\boldsymbol{w}, \mathcal{K}, \mathcal{A}}$ and obtain the same result. Altogether, this yields Alg. 3.

Algorithm 3 $(\epsilon, \delta, \theta)$-Certified Uniformity with Inverse Hessian Estimator

---

Require: Dataset $\mathcal{D}$; protected set $\mathcal{D}_p$; pretrained model $\boldsymbol{w}^* = \mathcal{A}(\mathcal{D})$; privacy budgets ϵ and δ; uniformity-utility tradeoff coefficient θ; estimator concentration b; sample size n; local convex coefficient λ; norm upper bound C; cumulative Hessian upper bound H; individual Hessian minimum eigenvalue upper bound $\zeta_{\min}$; gradient norm upper bound G; bound looseness probability ρ.

$\boldsymbol{P}_{0,\lambda}^{(0)} \leftarrow \nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$
for $j = 1, ..., b$ **do**
    for $t = 1, ..., n$ **do**
        Sample X_t from $\mathcal{D}_p$ uniformly with probability θ or,
        sample X_t from $\mathcal{D}_r$ uniformly with probability $1 - \theta$.
        if $X_t \sim \mathcal{D}_p$ **then**
            $\boldsymbol{H}_{t,\lambda}^{(j)} \leftarrow \nabla_{\boldsymbol{w}}^2 \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}^*, X_t) + \frac{\lambda \boldsymbol{I}}{2\theta}$.
        else if $X_t \sim \mathcal{D}_r$ **then**
            $\boldsymbol{H}_{t,\lambda}^{(j)} \leftarrow \nabla_{\boldsymbol{w}}^2 \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}^*, X_t) + \frac{\lambda \boldsymbol{I}}{2(1-\theta)}$.
        end if
        $\boldsymbol{P}_{t,\lambda}^{(j)} = \boldsymbol{P}_{0,\lambda}^{(0)} + (\boldsymbol{I} - \frac{\boldsymbol{H}_{t,\lambda}^{(j)}}{H}) \boldsymbol{P}_{t-1,\lambda}^{(j)}$.
    end for
end for
$\boldsymbol{P}_{n,\lambda} \leftarrow \frac{1}{b} \sum_{j=1}^{b} \boldsymbol{P}_{n,\lambda}^{(j)}$.
$\tilde{\boldsymbol{w}} \leftarrow \boldsymbol{w}^* - \frac{\boldsymbol{P}_{n,\lambda}}{H}$.
Compute Δ as the bound in Eq. (F.21).
$\sigma = \frac{\Delta}{\epsilon} \sqrt{2 \ln(1.25/\delta)}$
$\boldsymbol{w}^- \leftarrow \tilde{\boldsymbol{w}} + Y$ where $Y \sim \mathcal{N}(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$.
return $\boldsymbol{w}^-$.

G   Eliminating Hyperparameters in Certified Algorithms

Here, we summarize how to eliminate hyperparameters in Alg. 2:

- $\lambda_{\min}$ can be chosen as 0 by convex approximation, or it can be estimated using simple algorithms like Gershgorin's circle theorem or inverse power iteration.

- By Lemma H.1, λ can be chosen as $\theta P_{\mathcal{K}} + (1 - \theta) P_{\mathcal{A}}$

- Similarly, by Lemma H.1, H can be chosen as 2λ

- In practice, since $\frac{B}{\lambda + \lambda_{\min}}$ offers a bound on $\hat{\kappa}_l$ by Lemma H.6, $\zeta_{\min}$ can be chosen as $\lambda + \lambda_{\min}$

- By Lemma H.7, n can be chosen as $n = 2 \frac{B}{\lambda + \lambda_{\min}} \ln(\frac{B}{\lambda + \lambda_{\min}} b)$

- G can be approximated by computing $||\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2$.

- θ can be chosen with Cor. 5.2 to satisfy a particular closeness to uniformity.

- In practice, we find that C can be chosen as 10, 20, or 100.

- b can be chosen to satisfy a particular concentration on the estimator, so we let it be free. However, one can set $b = 1$.

- Common heuristics for ε and δ are available in the differential privacy literature.

- In practice, following what is common in certified unlearning e.g. in (Zhang et al., 2024), the Lipchitz constants in Asm. 5.3 and Asm. 5.4 are treated as hyperparameters. However, in practice, they can all be set to 1.

H Proofs

H.1 Helpful Lemmas

Lemma H.1. *Given Asm. 5.3, the gradients $\nabla_{\boldsymbol{w},\mathcal{K}}$ and $\nabla_{\boldsymbol{w},\mathcal{A}}$ exist and are Lipschitz with constants $P_{\mathcal{K}}$ and $P_{\mathcal{A}}$, respectively. Furthermore, given Asm. 5.4, the Hessians $\boldsymbol{H}_{\boldsymbol{w},\mathcal{K}}$ and $\boldsymbol{H}_{\boldsymbol{w},\mathcal{A}}$ exist and are Lipschitz with constants $F_{\mathcal{K}}$ and $F_{\mathcal{A}}$, respectively.*

Proof.

$$||\nabla_{\boldsymbol{w}_1,\mathcal{K}} - \nabla_{\boldsymbol{w}_2,\mathcal{K}}||_2 = ||\sum_{i=1}^{|\mathcal{D}_p|} \nabla \ell_{\mathcal{K}}^{(i)}(\boldsymbol{w}_1, \mathcal{D}_p^{(i)}) - \sum_{i=1}^{|\mathcal{D}_p|} \nabla \ell_{\mathcal{K}}^{(i)}(\boldsymbol{w}_2, \mathcal{D}_p^{(i)})||_2 \tag{H.22}$$

$$\leq \sum_{i=1}^{|\mathcal{D}_p|} ||\nabla \ell_{\mathcal{K}}^{(i)}(\boldsymbol{w}_1, \mathcal{D}_p^{(i)}) - \nabla \ell_{\mathcal{K}}^{(i)}(\boldsymbol{w}_2, \mathcal{D}_p^{(i)})||_2, \text{ triangle inequality} \tag{H.23}$$

$$\leq \sum_{i=1}^{|\mathcal{D}_p|} \frac{P_{\mathcal{K}}}{|\mathcal{D}_p|} ||\boldsymbol{w}_1 - \boldsymbol{w}_2||_2, \text{Asm. 5.3} \tag{H.24}$$

$$= P_{\mathcal{K}}||\boldsymbol{w}_1 - \boldsymbol{w}_2||_2 \tag{H.25}$$

This follows similarly for $\nabla_{\boldsymbol{w},\mathcal{A}}$, $\boldsymbol{H}_{\boldsymbol{w},\mathcal{K}}$, and $\boldsymbol{H}_{\boldsymbol{w},\mathcal{A}}$. $\square$

Lemma H.2. *Given Asm. 5.3, for any dataset $\mathcal{D} \subset \mathcal{Z}^n$, $\mathcal{L}_{\mathcal{A}}$ satisfies:*

$$|\mathcal{L}_A(\boldsymbol{w}_1, \mathcal{D}) - \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}_2, \mathcal{D})| \leq \frac{P_{\mathcal{A}}}{2}||\boldsymbol{w}_1 - \boldsymbol{w}_2||_2^2 + ||\nabla_{\boldsymbol{w}_2,\mathcal{A}}||_2||\boldsymbol{w}_1 - \boldsymbol{w}_2||_2 \tag{H.26}$$

Proof. By the fundamental theorem of calculus, we have the path integral:

$$\int_0^1 \langle \nabla_{\boldsymbol{w}_2+t(\boldsymbol{w}_1-\boldsymbol{w}_2),\mathcal{A}}, \boldsymbol{w}_1 - \boldsymbol{w}_2 \rangle dt = \mathcal{L}_A(\boldsymbol{w}_1, \mathcal{D}) - \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}_2, \mathcal{D}) \tag{H.27}$$

We have that:

$$\int_0^1 \langle \nabla_{\boldsymbol{w}_2+t(\boldsymbol{w}_1-\boldsymbol{w}_2),\mathcal{A}}, \boldsymbol{w}_1 - \boldsymbol{w}_2 \rangle dt = \int_0^1 \langle \nabla_{\boldsymbol{w}_2,\mathcal{A}} - \nabla_{\boldsymbol{w}_2,\mathcal{A}} + \nabla_{\boldsymbol{w}_2+t(\boldsymbol{w}_1-\boldsymbol{w}_2),\mathcal{A}}, \boldsymbol{w}_1 - \boldsymbol{w}_2 \rangle dt \tag{H.28}$$

$$= \int_0^1 \langle \nabla_{\boldsymbol{w}_2,\mathcal{A}}, \boldsymbol{w}_1 - \boldsymbol{w}_2 \rangle dt \tag{H.29}$$

$$+ \int_0^1 \langle \nabla_{\boldsymbol{w}_2+t(\boldsymbol{w}_1-\boldsymbol{w}_2),\mathcal{A}} - \nabla_{\boldsymbol{w}_2,\mathcal{A}}, \boldsymbol{w}_1 - \boldsymbol{w}_2 \rangle dt \tag{H.30}$$

The first term can be bounded by Cauchy-Schwarz as:

$$\int_0^1 \langle \nabla_{\boldsymbol{w}_2,\mathcal{A}}, \boldsymbol{w}_1 - \boldsymbol{w}_2 \rangle dt \leq ||\nabla_{\boldsymbol{w}_2,\mathcal{A}}||_2||\boldsymbol{w}_1 - \boldsymbol{w}_2||_2 \tag{H.31}$$

and similarly the second term can be bounded by Cauchy-Schwarz as:

$$\int_0^1 \langle \nabla_{\boldsymbol{w_2}+t(\boldsymbol{w_1}-\boldsymbol{w_2}),\mathcal{A}} - \nabla_{\boldsymbol{w_2},\mathcal{A}}, \boldsymbol{w_1}-\boldsymbol{w_2}\rangle dt \leq \int_0^1 ||\nabla_{\boldsymbol{w_2}+t(\boldsymbol{w_1}-\boldsymbol{w_2}),\mathcal{A}} - \nabla_{\boldsymbol{w_2},\mathcal{A}}||_2 ||\boldsymbol{w_1}-\boldsymbol{w_2}||_2 dt \qquad \text{(H.32)}$$

$$\leq \int_0^1 P_{\mathcal{K}} t ||\boldsymbol{w_1}-\boldsymbol{w_2}||_2, \text{by Lemma H.1} \qquad \text{(H.33)}$$

$$\leq \frac{P_{\mathcal{A}}}{2}||\boldsymbol{w_1}-\boldsymbol{w_2}||_2 \qquad \text{(H.34)}$$

Incorporating these bounds into Eq. (H.30) and Eq. (H.27), upon applying the triangle inequality, yields:

$$|\mathcal{L}_A(\boldsymbol{w_1},\mathcal{D}) - \mathcal{L}_{\mathcal{A}}(\boldsymbol{w_2},\mathcal{D})| \leq \frac{P_{\mathcal{A}}}{2}||\boldsymbol{w_1}-\boldsymbol{w_2}||_2^2 + ||\nabla_{\boldsymbol{w_2},\mathcal{A}}||_2||\boldsymbol{w_1}-\boldsymbol{w_2}||_2 \qquad \text{(H.35)}$$

as desired.

$\square$

Lemma H.3. *Given Asm. 5.4, the Hessians $\boldsymbol{H_{w,\mathcal{K}}}$, $\boldsymbol{H_{w,\mathcal{A}}}$, and $\boldsymbol{H_{w,\mathcal{K},\mathcal{A}}}$ are symmetric.*

Proof. By Lemma H.1, the Hessians $\boldsymbol{H_{w,\mathcal{K}}}$ and $\boldsymbol{H_{w,\mathcal{A}}}$ are continuous, and thus $\boldsymbol{H_{w,\mathcal{K},\mathcal{A}}}$ is continuous by linearity. Hence, all second-order partial derivatives contained in the Hessians are continuous, so by Schwartz's theorem all Hessians are symmetric. Importantly, for e.g. $\boldsymbol{H_{w,\mathcal{K},\mathcal{A}}}$, $||\boldsymbol{H_{w,\mathcal{K},\mathcal{A}}}||_2 = \max_i |\lambda_i(\boldsymbol{H_{w,\mathcal{K},\mathcal{A}}})|$, where λ_i denotes the ith eigenvalue. $\square$

Lemma H.4. *(Corollary of Theorem A.1 in (Dwork et al., 2014)) Let $X \sim \mathcal{N}(\lambda, \sigma^2 \boldsymbol{I})$ and $\mathcal{Y} \sim \mathcal{N}(\lambda', \sigma^2 \boldsymbol{I})$. Suppose $||\lambda - \lambda'||_2 \leq \Delta$. Then for any $\delta > 0$, X and Y are (ε, δ)-indistinguishable if $\sigma \geq \frac{\Delta}{\varepsilon}\sqrt{2\ln(1.25/\delta)}$.*

Lemma H.5. *Suppose we have n i.i.d. data samples $(X_1, ..., X_n)$ drawn from D_f and D_r with probabilities θ and $1 - \theta$ respectively. For $t = 1, ..., n$, if $X_t \sim D_f$ let $\boldsymbol{H}_{t,\lambda} = \boldsymbol{H_{w^*,\mathcal{K},t}} + \frac{\lambda \boldsymbol{I}}{2\theta}$ and if $X_t \sim D_r$ let $\boldsymbol{H}_{t,\lambda} = \boldsymbol{H_{w^*,\mathcal{A},t}} + \frac{\lambda \boldsymbol{I}}{2(1-\theta)}$. Then, $\mathbb{E}[\boldsymbol{H}_{t,\lambda}] = \boldsymbol{H_{w^*,\mathcal{K},\mathcal{A}}} + \lambda \boldsymbol{I}$ i.e. $\boldsymbol{H}_{t,\lambda}$ is an unbiased estimator of our Hessian of interest.*

Proof. At time t, we have sample X_t s.t. $X_t \sim D_r$ or $X_t \sim D_f$. Note that $\mathbb{E}_{X_t \sim D_f}[\boldsymbol{H_{w^*,\mathcal{K},t}} + \frac{\lambda \boldsymbol{I}}{2\theta}] = \boldsymbol{H_{w^*,\mathcal{K}}} + \frac{\lambda \boldsymbol{I}}{2\theta}$, and likewise $\mathbb{E}_{X_t \sim D_r}[\boldsymbol{H_{w^*,\mathcal{A},t}} + \frac{\lambda \boldsymbol{I}}{2(1-\theta)}] = \boldsymbol{H_{w^*,\mathcal{A}}} + \frac{\lambda \boldsymbol{I}}{2(1-\theta)}$.

By the law of iterated expectation, we have that:

$$\mathbb{E}[\boldsymbol{H}_{t,\lambda}] = \mathbb{E}[\boldsymbol{H}_{t,\lambda}|X_t \sim \mathcal{D}_p]\Pr(X_t \sim \mathcal{D}_p) + \mathbb{E}[\boldsymbol{H}_{t,\lambda}|X_t \sim \mathcal{D}_r]\Pr(X_t \sim \mathcal{D}_r)$$

$$= \theta\mathbb{E}_{X_t \sim D_f}[\boldsymbol{H_{w^*,\mathcal{K},t}} + \frac{\lambda \boldsymbol{I}}{2\theta})] + (1-\theta)\mathbb{E}_{X_t \sim \mathcal{D}_r}[\boldsymbol{H_{w^*,\mathcal{A},t}} + \frac{\lambda \boldsymbol{I}}{2(1-\theta)}]$$

$$= \theta(\boldsymbol{H_{w^*,\mathcal{K}}} + \frac{\lambda \boldsymbol{I}}{2\theta}) + (1-\theta)(\boldsymbol{H_{w^*,\mathcal{A}}} + \frac{\lambda \boldsymbol{I}}{2(1-\theta)})$$

$$= \theta\boldsymbol{H_{w^*,\mathcal{K}}} + (1-\theta)\boldsymbol{H_{w^*,\mathcal{A}}} + \lambda \boldsymbol{I}$$

$$= \boldsymbol{H_{w^*,\mathcal{K},\mathcal{A}}} + \lambda \boldsymbol{I}$$

as desired. $\square$

Lemma H.6. *Suppose Assumptions 5.3 and 5.4 hold. Let local condition number $\hat{\kappa}_l$ and maximum local condition number $\hat{\kappa}_l^{\max}$ correspond to the definitions of Agarwal et al. (2016) with respect to the Hessian of the loss of $\mathcal{M}_\theta$ after local convex approximation. Then, $\hat{\kappa}_l \leq \frac{B}{\lambda + \lambda_{\min}}$ and $\hat{\kappa}_l^{\max} \leq \frac{B}{\zeta_{\min}}$ where $B = \max\{\frac{\theta P_{\mathcal{K}}+\lambda}{|\mathcal{D}_p|}, \frac{(1-\theta)P_{\mathcal{A}}+\lambda}{|\mathcal{D}_r|}\}$ and where $\zeta_{\min} \leq \min_i \lambda_{\min}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}^{(i)}(\boldsymbol{w}, \mathcal{D}^{(i)}))$.*

Proof. By our definition of $\mathcal{M}_\theta$ and our local convex approximation technique, we have that:

$$M_\theta(D) = \arg\min_{\boldsymbol{w}\in\mathcal{W}} \ \theta \sum_{i=1}^{|\mathcal{D}_p|} \ell_\mathcal{K}^{(i)}(\boldsymbol{w}, \mathcal{D}_p^{(i)}) + (1-\theta)\sum_{i=1}^{|\mathcal{D}_r|} \ell_\mathcal{A}^{(i)}(\boldsymbol{w}, \mathcal{D}_r^{(i)}) + \frac{\lambda}{2}||\boldsymbol{w}||_2^2 \tag{H.36}$$

$$= \arg\min_{\boldsymbol{w}\in\mathcal{W}} \ \sum_{i=1}^{|D_f|} (\theta\ell_\mathcal{K}^{(i)}(\boldsymbol{w}, \mathcal{D}_p^{(i)}) + \frac{\lambda}{2|\mathcal{D}_p|}||\boldsymbol{w}||_2^2) + \sum_{i=1}^{|D_r|}(\theta\ell_\mathcal{A}^{(i)}(\boldsymbol{w}, \mathcal{D}_r^{(i)}) + \frac{\lambda}{2|\mathcal{D}_r|}||\boldsymbol{w}||_2^2) \tag{H.37}$$

$$= \arg\min_{\boldsymbol{w}\in\mathcal{W}} \ \sum_{i=1}^{|D|} \tilde{\ell}_{\mathcal{K},\mathcal{A}}^{(i)}(\boldsymbol{w}, \mathcal{D}^{(i)}) \tag{H.38}$$

where

$$\tilde{\ell}_{\mathcal{K},\mathcal{A}}^{(i)}(\boldsymbol{w}, \mathcal{D}^{(i)}) = \begin{cases} \theta\ell_\mathcal{K}^{(i)}(\boldsymbol{w}, \mathcal{D}_p^{(i)}) + \frac{\lambda}{2|\mathcal{D}_p|}||\boldsymbol{w}||_2^2, \ 1 \le i \le |\mathcal{D}_p| \\ (1-\theta)\ell_\mathcal{K}^{(i-|\mathcal{D}_p|)}(\boldsymbol{w}, \mathcal{D}_r^{(i-|\mathcal{D}_p|)}) + \frac{\lambda}{2|\mathcal{D}_r|}||\boldsymbol{w}||_2^2, \ |\mathcal{D}_p|+1 \le i \le |\mathcal{D}| \end{cases} \tag{H.39}$$

By the definitions provided in Agarwal et al. (2016), we have:

$$\hat{\kappa}_l = \max_{\boldsymbol{w}\in\mathcal{W}} \frac{\max_i \lambda_{\max}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}(\boldsymbol{w}, \mathcal{D}^{(i)}))}{\lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})} \tag{H.40}$$

and

$$\hat{\kappa}_l^{\max} = \max_{\boldsymbol{w}\in\mathcal{W}} \frac{\max_i \lambda_{\max}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}(\boldsymbol{w}, \mathcal{D}^{(i)}))}{\min_i \lambda_{\min}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}(\boldsymbol{w}, \mathcal{D}^{(i)}))} \tag{H.41}$$

We then have that, for any i,

$$\lambda_{\max}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}^{(i)}) \le ||\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}^{(i)}||_2, \text{ by Lemma H.3} \tag{H.42}$$

$$= \max\{||\theta\nabla_{\boldsymbol{w}}^2\ell_\mathcal{K}^{(i)} + \frac{\lambda\boldsymbol{I}}{|\mathcal{D}_p|}||_2, ||(1-\theta)\nabla_{\boldsymbol{w}}^2\ell_\mathcal{A}^{(i)} + \frac{\lambda\boldsymbol{I}}{|\mathcal{D}_r|}||_2\} \tag{H.43}$$

$$\tag{H.44}$$

Furthermore, by Asm. 5.3 and the triangle inequality:

$$||\theta\nabla_{\boldsymbol{w}}^2\ell_\mathcal{K}^{(i)} + \frac{\lambda\boldsymbol{I}}{|\mathcal{D}_p|}||_2 \le \frac{\theta P_\mathcal{K} + \lambda}{|\mathcal{D}_p|} \tag{H.45}$$

and

$$||(1-\theta)\nabla_{\boldsymbol{w}}^2\ell_\mathcal{A}^{(i)} + \frac{\lambda\boldsymbol{I}}{|\mathcal{D}_r|}||_2 \le \frac{(1-\theta)P_\mathcal{A} + \lambda}{|\mathcal{D}_r|} \tag{H.46}$$

Taking max over all i, we obtain that

$$\max_i \lambda_{\max}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K},\mathcal{A}}(\boldsymbol{w}, \mathcal{D}^{(i)})) \le \max\{\frac{\theta P_\mathcal{K} + \lambda}{|\mathcal{D}_p|}, \frac{(1-\theta)P_\mathcal{A} + \lambda}{|\mathcal{D}_r|}\} \tag{H.47}$$

which we denote by B.

Then, we obtain that $\hat{\kappa}_l \leq \frac{B}{\lambda + \lambda_{\min}}$ and $\hat{\kappa}_l^{\max} \leq \frac{B}{\zeta_{\min}}$ as desired, since

$\square$

Lemma H.7. *(Lemma 3.6 adapted from Agarwal et al. (2016)) Suppose Asm. 5.3 and Asm. 5.4 hold. Consider the estimator $\frac{\tilde{H}_{n,\lambda}^{-1}}{H}$ in Theorem F.4. Let b be the number of inverse Hessian estimators we obtain. Suppose $n \geq 2\frac{B}{\lambda + \lambda_{\min}} \ln(\frac{B}{\lambda + \lambda_{\min}} b)$, where $B = \max\{\frac{\theta P_\mathcal{K} + \lambda}{|\mathcal{D}_p|}, \frac{(1-\theta) P_\mathcal{A} + \lambda}{|\mathcal{D}_r|}\}$. Then, we have that:*

$$\Pr[||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} + \lambda \boldsymbol{I})^{-1} - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}||_2 \leq 16 \frac{B}{\zeta_{\min}} \sqrt{\frac{\ln(\frac{z}{\rho})}{b}} + \frac{1}{16}] \geq 1 - \rho \qquad (\text{H.48})$$

where $\zeta_{\min} \leq \min_i \lambda_{\min}(\nabla_{\boldsymbol{w}}^2 \tilde{\ell}_{\mathcal{K}, \mathcal{A}}^{(i)}(\boldsymbol{w}, \mathcal{D}^{(i)}))$.

Proof. Note that $b = S_1$ in our setting. In our setting, following the subsequent steps of the proof in Agarwal et al. (2016) after plugging in the bounds in Lemma H.6 in place of $\hat{\kappa}_l, \hat{\kappa}_l^{\max}$, noting that we choose $n = S_2 \geq 2\frac{B}{\lambda + \lambda_{\min}} \ln(\frac{B}{\lambda + \lambda_{\min}} b)$, we obtain the exact same result for the Neumann series bound of $\frac{1}{16}$. Using the fact that $\frac{B}{\zeta_{\min}}$ is an upper bound on $\hat{\kappa}_l^{\max}$ by Lemma H.6, the rest of the proof follows similarly. $\square$

Lemma H.8. *(Proposition 2.1 in Dwork et al. (2014)) Let $\mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \to R$ be a randomized algorithm that is (ε, δ)-differentially private. Let $f : R \to R'$ be an arbitrary mapping. Then, $f \circ \mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \to R'$ is (ε, δ)-differentially private.*

Note that, in the proof of Lemma H.8, one proves this fact for deterministic mappings, so this holds for both randomized and deterministic f.

Lemma H.9. *Consider the mapping $\mathcal{J} : \mathcal{W} \to R$, and suppose $\mathcal{G} : \mathcal{Z}^n \times \mathcal{Z}^n \times \mathcal{W} \to \mathcal{W}$ satisfies Def. 4.4. Then, $\forall C \subset R$:*

$$\Pr(\mathcal{J}(\mathcal{G}(\mathcal{D}, \mathcal{D}_p, \mathcal{A}(\mathcal{D}))) \in C) \leq e^\varepsilon \Pr(\mathcal{J}(\mathcal{M}_\theta(\mathcal{D})) \in C) + \delta \qquad (\text{H.49})$$
$$\Pr(\mathcal{J}(\mathcal{M}_\theta(\mathcal{D})) \in C) \leq e^\varepsilon \Pr(\mathcal{J}(\mathcal{G}(\mathcal{D}, \mathcal{D}_p, \mathcal{A}(\mathcal{D}))) \in C) + \delta \qquad (\text{H.50})$$

Proof. Immediate from Lemma H.8. $\square$

H.2 Proof of Proposition H.10

We seek to prove the following statement:

Proposition H.10. *Suppose we have a hypothesis space $\mathcal{H}_\mathcal{W}$ consisting of functions where the ultimate layer is an affine transformation and the outputs are passed through a softmax. Let $\mathcal{K}$ be a uniform learner. Then, $f_{\mathcal{K}(\mathcal{D})} \in \mathcal{H}_\mathcal{W} \; \forall \mathcal{D} \subset \mathcal{Z}^n$.*

Proof. Fix a dataset $\mathcal{D} \subset \mathcal{Z}^n$.

Suppose we have an $K-$layer function $f_{\boldsymbol{w}} : \mathbb{R}^d \to \mathbb{R}^o$ parameterized by $\boldsymbol{w} \in \mathcal{W}$ of the form $f(\boldsymbol{x}) = L_1 \circ ... \circ L_{K-1} \circ L_K$ where $L_{K-1}(\boldsymbol{x}) = \boldsymbol{W}_{K-1}^T \boldsymbol{x} + \boldsymbol{b}_{K-1}$ and $L_K(\boldsymbol{x}) = \text{softmax}(\boldsymbol{x})$, i.e. $L_K(\boldsymbol{x})_i = \frac{e^{x_i}}{\sum_{j=1}^{|\mathcal{Y}|} e^{x_j}}$.

Thus, $f_{\boldsymbol{w}} \in \mathcal{H}_\mathcal{W}$. Then, let $\boldsymbol{W}_{K-1} = \boldsymbol{0}$ and $\boldsymbol{b}_{K-1} = \boldsymbol{0}$.

Fix $\boldsymbol{z} \in \mathcal{D}$. This yields, for $j = 1, ..., |\mathcal{Y}|$, $f(\boldsymbol{z})_j = \frac{e^0}{\sum_{j=1}^{|\mathcal{Y}|} e^0} = \frac{e^0}{|\mathcal{Y}| e^0} = \frac{1}{|\mathcal{Y}|}$. Hence, since $\boldsymbol{z}$ was arbitrary,

$f_{\boldsymbol{w}}(\boldsymbol{z}) = (\underbrace{\frac{1}{|\mathcal{Y}|}, ..., \frac{1}{|\mathcal{Y}|}}_{|\mathcal{Y}| \text{ times}}) \; \forall \boldsymbol{z} \in \mathcal{D}$. Since $\mathcal{D}$ was arbitrary, by definition of a uniform learner over $\mathcal{D}$, $f_{\mathcal{K}(\mathcal{D})} \in$

$\mathcal{H}_\mathcal{W} \; \forall \mathcal{D} \subset \mathcal{Z}^n$ as desired.

$\square$

H.3 Proof of Proposition H.11

We seek to prove the following statement:

Proposition H.11. *Let $\theta \in (0,1)$. Fix $\mathcal{D} \subset \mathcal{Z}^n$ and consider the* protected set *$\mathcal{D}_p \subset \mathcal{D}$ and the retain set $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_p$. Then, if is a global minimizer, it is globally Pareto optimal with respect to $\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p)$ and $\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)$. Similarly, if $\mathcal{M}_\theta(\mathcal{D})$ is a local minimizer, it is locally Pareto optimal.*

We use the following definition of global Pareto optimality:

Definition H.12. *(Chapter 1 of Pardalos et al. (2017)) Suppose we have a multiobjective optimization problem $\min \boldsymbol{f}(\boldsymbol{x})$ s.t. $\boldsymbol{x} \in A$, where $\boldsymbol{f}(\boldsymbol{x}) = (f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), ..., f_m(\boldsymbol{x}))$. $\boldsymbol{x}^* \in A$ with $\boldsymbol{f}(\boldsymbol{x}^*)$ is called globally Pareto optimal if and only if there exists no $\boldsymbol{x} \in A$ such that $f_i(\boldsymbol{x}) \le f_i(\boldsymbol{x}^*)$ for all $i = 1, 2, ..., m$ and $f_j(\boldsymbol{x}) < f_j(\boldsymbol{x}^*)$ for at least one $j \in \{1, ..., m\}$.*

We can then prove the statement:

Proof. Let $\theta \in (0,1)$. Fix $\mathcal{D} \subset \mathcal{Z}^n$, $\mathcal{D}_p \subset \mathcal{D}$, and $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_p$.

Suppose, for the sake of contradiction, that $\tilde{\boldsymbol{w}}^* = \mathcal{M}_\theta(\mathcal{D}) = \arg\min_{\boldsymbol{w}} \theta \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)$, a global minimizer, is not globally Pareto optimal with respect to $\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p)$ and $\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)$. Then, exists $\boldsymbol{w}'$ s.t. $\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}', \mathcal{D}_p) \le \mathcal{L}_{\mathcal{K}}(\tilde{\boldsymbol{w}}^*, \mathcal{D}_p)$ and $\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}', \mathcal{D}_r) \le \mathcal{L}_{\mathcal{A}}(\tilde{\boldsymbol{w}}^*, \mathcal{D}_r)$, with at least one of these inequalities being strict.

Then, since $\theta \in (0,1)$ and $(1-\theta) \in (0,1)$, we have that $\theta \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}', \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}', \mathcal{D}_r) < \theta \mathcal{L}_{\mathcal{K}}(\tilde{\boldsymbol{w}}^*, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\tilde{\boldsymbol{w}}^*, \mathcal{D}_r)$, contradicting optimality of $\tilde{\boldsymbol{w}}^*$. As such, $\mathcal{M}_\theta(\mathcal{D})$ is globally Pareto optimal respect to $\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p)$ and $\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)$ as desired.

This holds similarly for a local minimizer $\tilde{\boldsymbol{w}}^*$, where Pareto optimality similarly holds only locally in a neighborhood around the minima.

$\square$

H.4 Proof of Theorem F.1

Proof. The proof follows similarly to Lemma 10 in Sekhari et al. (2021); for completeness, we adapt their proof to our setting.

Let $\boldsymbol{w}^* := A(D), \boldsymbol{w}^- := \mathcal{G}(\mathcal{D}, \mathcal{D}_p, \boldsymbol{w}^*), \tilde{\boldsymbol{w}} := \mathcal{F}(\mathcal{D}, \mathcal{D}_p, \boldsymbol{w}^*)$. Departing from the notation of the theorem for clarity, let $\hat{\boldsymbol{w}}^* := \mathcal{M}_\theta(\mathcal{D}), \hat{\boldsymbol{w}}^- := \mathcal{G}(\mathcal{D}, \emptyset, \hat{\boldsymbol{w}}^*), \hat{\tilde{\boldsymbol{w}}} := \mathcal{F}(\mathcal{D}, \emptyset, \hat{\boldsymbol{w}}^*)$.

Note that $\hat{\tilde{\boldsymbol{w}}} = \hat{\boldsymbol{w}}^*$. We then have that $||\tilde{\boldsymbol{w}} - \hat{\tilde{\boldsymbol{w}}}||_2 = ||\tilde{\boldsymbol{w}} - \hat{\boldsymbol{w}}^*||_2 \le \Delta$, by definition of Δ.

By definition of $\mathcal{G}$, we have that $\boldsymbol{w}^- = \tilde{\boldsymbol{w}} + Y$ and $\hat{\boldsymbol{w}}^- = \hat{\tilde{\boldsymbol{w}}} + Y$, where $Y \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I})$ s.t. $\sigma \ge \frac{\Delta}{\varepsilon}\sqrt{2\ln(1.25/\delta)}$.

As such, $\boldsymbol{w}^- = \mathcal{N}(\tilde{\boldsymbol{w}}, \sigma^2 \boldsymbol{I})$ and $\hat{\boldsymbol{w}}^- \sim \mathcal{N}(\hat{\tilde{\boldsymbol{w}}}, \sigma^2 \boldsymbol{I})$.

Thus, by Lemma H.4, $\boldsymbol{w}^-, \hat{\boldsymbol{w}}^-$ are (ε, δ)-indistinguishable. In particular, $\mathcal{G}(\mathcal{D}, \mathcal{D}_p, \mathcal{A}(\mathcal{D}))$ and $\mathcal{M}_\theta(\mathcal{D})$ are (ε, δ)-indistinguishable, as desired.

$\square$

H.5 Proof of Proposition F.2

Proof. By the same token as Lemma 3.3 in (Zhang et al., 2024), we have that:

$$||\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^*||_2 \le ||\boldsymbol{H}^{-1}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}||_2 \int_0^1 ||\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}} - \boldsymbol{H}_{\boldsymbol{w}^* + t(\tilde{\boldsymbol{w}}^* - \boldsymbol{w}^*), \mathcal{K}, \mathcal{A}}||_2 ||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2 dt \qquad \text{(H.51)}$$

.

Let $\boldsymbol{w}' = \boldsymbol{w}^* + t(\tilde{\boldsymbol{w}} - \boldsymbol{w}^*)$. We have that $||\boldsymbol{w}^* - \boldsymbol{w}'||_2 = ||\boldsymbol{w}^* - \boldsymbol{w}^* + t(\tilde{\boldsymbol{w}}^* - \boldsymbol{w}^*)||_2 = t||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2$.

Furthermore, by linearity of $\boldsymbol{H_w}$ and the triangle inequality, we have that:

$$||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \boldsymbol{H}_{\boldsymbol{w}',\mathcal{K},\mathcal{A}}||_2 = ||\theta \boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K}} + (1-\theta)\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{A}} - \theta \boldsymbol{H}_{\boldsymbol{w}',\mathcal{K}} - (1-\theta)\boldsymbol{H}_{\boldsymbol{w}',\mathcal{A}}||_2 \tag{H.52}$$

$$\leq \theta||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K}} - \theta \boldsymbol{H}_{\boldsymbol{w}',\mathcal{K}}||_2 + (1-\theta)||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{A}} - \boldsymbol{H}_{\boldsymbol{w}',\mathcal{A}}||_2 \tag{H.53}$$

$$= \theta F_{\mathcal{K}}||\boldsymbol{w}^* - \boldsymbol{w}'||_2 + (1-\theta)F_{\mathcal{A}}||\boldsymbol{w}^* - \boldsymbol{w}'||_2, \text{ by Lemma H.1} \tag{H.54}$$

$$= \theta t F_{\mathcal{K}}||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2 + (1-\theta)t F_{\mathcal{A}}||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2, \tag{H.55}$$

This yields that:

$$||\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^*||_2 \leq ||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}}^{-1}||_2 \int_0^1 ||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \boldsymbol{H}_{\boldsymbol{w}^*+t(\tilde{\boldsymbol{w}}^*-\boldsymbol{w}^*),\mathcal{K},\mathcal{A}}||_2 ||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2 dt \tag{H.56}$$

$$\leq ||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}}^{-1}||_2 \int_0^1 (\theta t F_{\mathcal{K}} + (1-\theta)t F_{\mathcal{A}})||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2^2 \tag{H.57}$$

$$= \frac{\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}}}{2}||\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}}^{-1}||_2 ||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^*||_2^2 \tag{H.58}$$

as desired.

$\square$

H.6 Proof of Proposition F.3

Proof. See the proof of theorem 3.4 in (Zhang et al., 2024), noting that in our setting $M = F = \theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}}$ by Eq. (H.55). $\square$

H.7 Proof of Theorem F.4

Proof. First, we have that:

$$\mathbb{E}[\tilde{\boldsymbol{H}}_{t,\lambda}^{-1}] = \mathbb{E}[\boldsymbol{I} + \tilde{\boldsymbol{H}}_{t-1,\lambda}^{-1} - \frac{1}{J}\boldsymbol{H}_{t,\lambda}\tilde{\boldsymbol{H}}_{t-1,\lambda}^{-1}], \text{ by definition} \tag{H.59}$$

$$= \boldsymbol{I} + \mathbb{E}[\tilde{\boldsymbol{H}}_{t-1,\lambda}^{-1}] - \frac{1}{J}\mathbb{E}[\boldsymbol{H}_{t,\lambda}\tilde{\boldsymbol{H}}_{t-1,\lambda}], \text{ linearity of expectation} \tag{H.60}$$

$$= \boldsymbol{I} + \mathbb{E}[\tilde{\boldsymbol{H}}_{t-1,\lambda}^{-1}] - \frac{1}{J}\mathbb{E}[\boldsymbol{H}_{t,\lambda}]\mathbb{E}[\tilde{\boldsymbol{H}}_{t-1,\lambda}], \text{i.i.d. samples} \tag{H.61}$$

$$= \boldsymbol{I} + \mathbb{E}[\tilde{\boldsymbol{H}}_{t-1,\lambda}^{-1}] - \frac{\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I}}{J}\mathbb{E}[\tilde{\boldsymbol{H}}_{t-1,\lambda}], \text{ by Lemma H.5} \tag{H.62}$$

Denote $\boldsymbol{H}_* := \boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}}$ and $\boldsymbol{E}_t := \mathbb{E}[\tilde{\boldsymbol{H}}_{t,\lambda}^{-1}]$. We thus have that:

$$\boldsymbol{E}_t = \boldsymbol{I} + \boldsymbol{E}_{t-1} - \frac{\boldsymbol{H}_*}{J}\boldsymbol{E}_{t-1} \tag{H.63}$$

$$= \boldsymbol{I} + \boldsymbol{E}_{t-1}(\boldsymbol{I} - \frac{\boldsymbol{H}_*}{J}) \tag{H.64}$$

$$= \boldsymbol{I} + (\boldsymbol{I} - \boldsymbol{M})\boldsymbol{E}_{t-1}, \text{ letting } \boldsymbol{M} := \frac{\boldsymbol{H}_*}{J} \tag{H.65}$$

We then know that, by assumption, $\lambda > ||\boldsymbol{H}_*||_2$, where $\boldsymbol{H}_*$ is a symmetric Hessian by Lemma H.3; as such, $\boldsymbol{H}_* + \lambda \boldsymbol{I}$ is positive definite and has all positive eigenvalues. We also know that $||\boldsymbol{H}_*||_2 < J \implies ||\boldsymbol{M}||_2 < 1$, so we have that $0 < \lambda_i(\boldsymbol{M}) < 1$ for all eigenvalues λ_i. Furthermore, $\boldsymbol{I} - \boldsymbol{M}$ has eigenvalues $1 - \lambda_i(\boldsymbol{M})$, so we have that $0 < \lambda_i(\boldsymbol{I} - \boldsymbol{M}) < 1$, so $||\boldsymbol{I} - \boldsymbol{M}||_2 < 1$, since $\boldsymbol{I} - \boldsymbol{M}$ is symmetric. Since $\boldsymbol{I} - \boldsymbol{M}$ has spectral radius less than 1, the Neumann series $\sum_{k=0}^\infty (\boldsymbol{I} - \boldsymbol{M})^k$ converges. (Mayer, 1985). Thus, the Neumann series is Cauchy.

Fix $\varepsilon > 0$. Let $s_n = \sum_{k=0}^n (\boldsymbol{I} - \boldsymbol{M})^k$. We know that $\exists N \in \mathbb{N}$ s.t. $m > n \geq N \implies ||s_m - s_n||_2 = ||\sum_{k=n+1}^m (\boldsymbol{I} - \boldsymbol{M})^k||_2 < \varepsilon$. For $m > n \geq N$, we have that $||\boldsymbol{E}_m - \boldsymbol{E}_n||_2 = ||\sum_{k=n+1}^m (\boldsymbol{I} - \boldsymbol{M})^k||_2 < \varepsilon$. As such, $\{\boldsymbol{E}_n\}$ is Cauchy; since it is real, it converges. As such, $\boldsymbol{E}_\infty = \lim_{t \to \infty} \boldsymbol{E}_n$ exists.

Taking limits on both sides, we then have:

$$\mathbb{E}[\tilde{\boldsymbol{H}}_{\infty,\lambda}^{-1}] = \boldsymbol{I} + \mathbb{E}[\tilde{\boldsymbol{H}}_{\infty,\lambda}^{-1}] + \frac{\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I}}{J} \mathbb{E}[\tilde{\boldsymbol{H}}_{\infty,\lambda}] \tag{H.66}$$

$$\iff \mathbb{E}[\frac{\tilde{\boldsymbol{H}}_{\infty,\lambda}^{-1}}{J}] = (\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I})^{-1} \tag{H.67}$$

rearranging using linearity of expectation and noting that λ was chosen such that $\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I}$ is invertible, as desired. $\square$

H.8 Proof of Theorem F.5

This follows similarly to theorem 3.6 and proposition 4.1 in Zhang et al. (2024), noting that we apply Lemma H.7 instead of applying lemma 3.6 from Agarwal et al. (2016). Furthermore, note that $L = \theta P_\mathcal{K} + (1 - \theta) P_\mathcal{A}$ in our setting. For completeness, we provide the full proof below.

Proof.

$$\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^* = \boldsymbol{w}^* - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H} \nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \tilde{\boldsymbol{w}}^* \tag{H.68}$$

$$= \boldsymbol{w}^* - \tilde{\boldsymbol{w}}^* - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}) - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H} \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}} \tag{H.69}$$

By the triangle inequality, this yields:

$$||\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^*||_2 \leq ||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^* - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 + ||\frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H} \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}||_2 \tag{H.70}$$

The first term in Eq. (H.70) can be bounded by the triangle inequality as:

$$||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^* - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \tag{H.71}$$

$$= ||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^* - ((\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I})^{-1} + \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H} - (\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I})^{-1})(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \tag{H.72}$$

$$\leq ||\boldsymbol{w}^* - \tilde{\boldsymbol{w}}^* - (\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I})^{-1}(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \tag{H.73}$$

$$+ ||((\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I})^{-1} - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H})(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \tag{H.74}$$

In the setting of Prop. F.3, we have that $\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^* = \boldsymbol{w}^* - \tilde{\boldsymbol{w}}^* - (\boldsymbol{H}_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} + \lambda \boldsymbol{I})^{-1}(\nabla_{\boldsymbol{w}^*,\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})$. Hence, by Prop. F.3, we have that:

$$||\boldsymbol{w^*} - \tilde{\boldsymbol{w}}^* - (\boldsymbol{H}_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1}(\nabla_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \tag{H.75}$$

$$\leq \frac{2C((\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})C + \lambda)}{\lambda + \lambda_{\min}} \tag{H.76}$$

Furthermore, we have:

$$||((\boldsymbol{H}_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1} - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H})(\nabla_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \tag{H.77}$$

$$\leq ||((\boldsymbol{H}_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1} - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H})||_2||(\nabla_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2, \text{property of op norm} \tag{H.78}$$

$$\leq (16\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln\frac{z}{\rho}}{b}} + \frac{1}{16})||(\nabla_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2, \text{ Lemma H.7} \tag{H.79}$$

$$\leq (16\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln\frac{z}{\rho}}{b}} + \frac{1}{16})2C(\theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}}, \text{ Lemma H.1} \tag{H.80}$$

with probability at least $1 - \rho$. Incorporating this into equation Eq. (H.74), we have that:

$$||\boldsymbol{w^*} - \tilde{\boldsymbol{w}}^* - \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}(\nabla_{\boldsymbol{w^*},\mathcal{K},\mathcal{A}} - \nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}})||_2 \leq \frac{2C((\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})C + \lambda)}{\lambda + \lambda_{\min}} \tag{H.81}$$

$$+ (32\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln\frac{z}{\rho}}{b}} + \frac{1}{8})C(\theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}}) \tag{H.82}$$

It then suffices to bound the second term in Eq. (H.70). We have that:

$$\left\|\frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}\right\|_2 = \left\|\left[(\boldsymbol{H}_{w^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1} - (\boldsymbol{H}_{w^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1}\right.\right.$$
$$\left.\left. + \frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H}\right]\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}\right\|_2 \tag{H.83}$$

$$= \left\|(\boldsymbol{H}_{w^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1}\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}\right.$$
$$\left. + \left(\frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H} - (\boldsymbol{H}_{w^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1}\right)\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}\right\|_2 \tag{H.84}$$

$$\leq \|(\boldsymbol{H}_{w^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1}\|_2\|\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}\|_2$$
$$+ \left\|\frac{\tilde{\boldsymbol{H}}_{n,\lambda}^{-1}}{H} - (\boldsymbol{H}_{w^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})^{-1}\right\|_2\|\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}\|_2 \tag{H.85}$$

$$\leq \frac{G}{\lambda + \lambda_{\min}} + \left(16\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln(z/\rho)}{b}} + \frac{1}{16}\right)G. \tag{H.86}$$

by definition of λ, Lemma H.7, and that $||\nabla_{\tilde{\boldsymbol{w}}^*,\mathcal{K},\mathcal{A}}||_2 \leq G$.

Incorporating the above into Eq. (H.70), this yields that:

$$\|\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^*\|_2 \leq \frac{2C\big((\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})\,C + \lambda\big)}{\lambda + \lambda_{\min}} \tag{H.87}$$

$$+ \Big(32\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln(z/\rho)}{b}} + \tfrac{1}{8}\Big)C\Big(\theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}} + \frac{G}{\lambda + \lambda_{\min}}$$

$$+ \Big(16\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln(z/\rho)}{b}} + \tfrac{1}{16}\Big)G\Big) \tag{H.88}$$

$$= \frac{2C\big((\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})\,C + \mu\big)\,+\,G}{\mu + \mu_{\min}} \tag{H.89}$$

$$+ \Big(16\frac{B}{\zeta_{\min}}\sqrt{\frac{\ln(z/\rho)}{b}} + \tfrac{1}{16}\Big)\big(2C(\theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}} + G)\big).$$

as desired.

$\square$

H.9 Proof of Proposition 5.1

Proof. Fix any sampled $\mathcal{D}$. Since $\mathcal{M}_\theta(\mathcal{D})$ is taken to be the global risk minimizer, we have that:

$$\theta\mathcal{L}_{\mathcal{K}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r) \tag{H.90}$$

$$\leq \theta\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)\ \forall \boldsymbol{w} \in \mathcal{W} \tag{H.91}$$

subtracting $\frac{\lambda}{2}\|\boldsymbol{w}\|_2$ from both sides.

Let $\boldsymbol{w}_U$ be the parameter that results in a parameterized model $f_{\boldsymbol{w}_U}$ which outputs a uniform distribution; by Prop. H.10, such a parameter exists. We then have that:

$$\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}_U, \mathcal{D}_p) = \sum_{i=1}^{|\mathcal{D}_p|} D_{KL}(U[0, |\mathcal{Y}|]) \|U[0, |\mathcal{Y}|] = 0 \tag{H.92}$$

and

$$\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}_U, \mathcal{D}_r) = \sum_{i=1}^{|\mathcal{D}_r|} \mathbb{H}_{CE}(\boldsymbol{y}^{(i)}, U[0, |\mathcal{Y}|]) = -\sum_{i=1}^{|\mathcal{D}_r|}\sum_{j=1}^{|\mathcal{Y}|} \boldsymbol{y}_j^{(i)} \ln\frac{1}{|\mathcal{Y}|} = |\mathcal{D}_r|\ln|\mathcal{Y}| \tag{H.93}$$

where $\boldsymbol{y}$ is a one hot vector of length $\mathcal{Y}$ such that for $\boldsymbol{y}_j^{(i)}$, $j = 1, ..., |\mathcal{Y}|$,

$$\boldsymbol{y}_j^{(i)} = \begin{cases} 1 & \text{instance } i \text{ is labeled class } j \\ 0 & \text{instance } i \text{ is not labeled class } j \end{cases} \tag{H.94}$$

Incorporating the above into Eq. (H.91) yields:

$$\theta\mathcal{L}_{\mathcal{K}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r) \leq \theta\mathcal{L}_{\mathcal{K}}(\boldsymbol{w}_U, \mathcal{D}_p) + (1-\theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}_U, \mathcal{D}_r) \tag{H.95}$$

$$\leq \theta(0) + (1-\theta)|\mathcal{D}_r|\ln|\mathcal{Y}| \tag{H.96}$$

$$= |\mathcal{D}_r|(1-\theta)\ln|\mathcal{Y}| \tag{H.97}$$

This then yields that:

$$\mathcal{L}_{\mathcal{K}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_p) \leq \frac{1-\theta}{\theta}(|\mathcal{D}_r| \ln |\mathcal{Y}| - \mathcal{L}_{\mathcal{A}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r) \tag{H.98}$$

$$\leq \frac{1-\theta}{\theta}|\mathcal{D}_r| \ln |\mathcal{Y}| \tag{H.99}$$

since the cross entropy is nonnegative, yielding that $-\mathcal{L}_{\mathcal{A}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r) \leq 0$.

Then, we have:

$$||f_{\mathcal{M}_\theta(\mathcal{D})}(\mathcal{D}_p) - U[0, |\mathcal{Y}|]||_\infty \leq ||f_{\mathcal{M}_\theta(\mathcal{D})}(\mathcal{D}_p) - U[0, |\mathcal{Y}|]||_1 \tag{H.100}$$

$$\leq 2TV(f_{\mathcal{M}_\theta(\mathcal{D})}(\mathcal{D}_p) - U[0, |\mathcal{Y}|]) \tag{H.101}$$

$$\leq 2\sqrt{\frac{1}{2}D_{KL}(f_{\mathcal{M}_\theta(\mathcal{D})}, ||U[0, |\mathcal{Y}|])}, \text{ Pinsker's inequality (Pinsker, 1964)} \tag{H.102}$$

$$= \sqrt{2D_{KL}(f_{\mathcal{M}_\theta(\mathcal{D})}, ||U[0, |\mathcal{Y}|])} \tag{H.103}$$

$$= \sqrt{2\mathcal{L}_{\mathcal{K}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_p)} \tag{H.104}$$

$$\leq \sqrt{2|\mathcal{D}_r|(\frac{1-\theta}{\theta}) \ln |\mathcal{Y}|} \tag{H.105}$$

by the above bound on $\mathcal{L}_{\mathcal{K}}$, as desired. $\qquad\square$

H.10 Proof of Corollary 5.2

Proof. To have Eq. (5), by Prop. 5.1, it suffices to solve for θ in the bound obtained. This results in:

$$\sqrt{2(\frac{1-\theta}{\theta})|\mathcal{D}_r| \ln |\mathcal{Y}|} \leq \gamma \iff \frac{1-\theta}{\theta}|\mathcal{D}_r| \ln |\mathcal{Y} \leq \frac{\gamma^2}{2} \tag{H.106}$$

$$\iff \frac{|\mathcal{D}_r| \ln |\mathcal{Y}|}{\theta} - \frac{|\mathcal{D}_r| \ln |\mathcal{Y}|\theta}{\theta} \leq \frac{\gamma^2}{2} \tag{H.107}$$

$$\iff \frac{|\mathcal{D}_r| \ln |\mathcal{Y}|}{\theta} \leq \frac{\gamma^2}{2} + |\mathcal{D}_r| \ln |\mathcal{Y}| = \frac{\gamma^2 + 2|\mathcal{D}_r| \ln |\mathcal{Y}|}{2} \tag{H.108}$$

$$\iff \frac{\theta}{|\mathcal{D}_r| \ln |\mathcal{Y}|} \geq \frac{2}{\gamma^2 + 2|\mathcal{D}_r| \ln |\mathcal{Y}|} \tag{H.109}$$

$$\iff \theta \geq \frac{2|\mathcal{D}_r| \ln |\mathcal{Y}}{\gamma^2 + 2|\mathcal{D}_r| \ln |\mathcal{Y}|} \tag{H.110}$$

as desired.

$\qquad\square$

H.11 Proof of Theorem 5.5

First, before we prove Theorem 5.5, we note that we can use Lemma H.2 and that $||\boldsymbol{w}|| \leq 2$ to obtain a simple bound. Let:

$$|\alpha^* - \alpha(\theta)| = |\mathcal{L}_{\mathcal{A}}(\mathcal{A}(\mathcal{D}_r), \mathcal{D}_r) - \mathcal{L}_{\mathcal{A}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r)| \tag{H.111}$$

$$\leq \frac{P_{\mathcal{A}}}{2}||\mathcal{M}_\theta(\mathcal{D}) - \mathcal{A}(\mathcal{D}_r)||_2^2 + ||\nabla_{\mathcal{A}(\mathcal{D}_r), \mathcal{A}}||_2||\mathcal{M}_\theta(\mathcal{D}) - \mathcal{A}(\mathcal{D}_r)||_2 \tag{H.112}$$

$$\leq \frac{C^2 P_{\mathcal{A}}}{2} + \lambda C^2 \tag{H.113}$$

after applying the triangle inequality and rearranging the first order condition on $\mathcal{A}(\mathcal{D}_r)$.

However, this bound is vacuous and not tight; it does not incorporate any information about θ or most of the constants that appear in Asm. 5.3 and Asm. 5.4. Given this, we seek to construct a tighter, non-vacuous bound. We first restate the proof without any asymptotic characterizations and concretely define how large λ must be:

Theorem H.13. *Suppose Assumptions 5.3 and 5.4 hold, and let $P_{\mathcal{K}}, P_{\mathcal{K}}, F_{\mathcal{K}}, F_{\mathcal{A}}$ be as defined in Assumptions 5.3 and 5.4. Let $\alpha^* := \mathcal{L}_{\mathcal{A}}(\mathcal{A}(\mathcal{D}_r), \mathcal{D}_r)$ be the locally optimal (empirical) retain loss, achieved by $\mathcal{M}_\theta(\mathcal{D})$ when $\theta = 0$. Let $\alpha(\theta) := \mathcal{L}_{\mathcal{A}}(\mathcal{M}_\theta(\mathcal{D}), \mathcal{D}_r)$ be the locally optimal retain loss obtained by $\mathcal{M}_\theta(\mathcal{D})$ when $\theta \in (0, 1)$. Suppose all weights used throughout are bounded by $||\boldsymbol{w}||_2 \leq C$. Additionally, denote by $F := \theta M_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}}$ and $P := \theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}}$. Consider regularization coefficient $\lambda \geq P + 2\theta CF + 2\sqrt{\theta CF(P + 2\theta CF + 8P_{\mathcal{K}})}$. Then, we have the following bound:*

$$|\alpha^* - \alpha(\theta)| \leq \frac{P_{\mathcal{K}}}{2}\left(\frac{\lambda - P - \sqrt{(\lambda - P)^2 - 4\theta CF(2P_{\mathcal{K}} + \lambda)}}{2F}\right)^2 + \tag{H.114}$$

$$\lambda C\left(\frac{\lambda - P - \sqrt{(\lambda - P)^2 - 4\theta CF(2P_{\mathcal{K}} + \lambda)}}{2F}\right). \tag{H.115}$$

Proof. First, when $\theta = 0$, we have that:

$$\boldsymbol{w}_{\alpha^*} := \arg\min_{\boldsymbol{w} \in \mathcal{W}, ||\boldsymbol{w}||_2 \leq C} \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r) + \frac{\lambda}{2}||\boldsymbol{w}||_2^2 \tag{H.116}$$

which yields the first order condition:

$$\nabla_{\boldsymbol{w}_{\alpha^*}, \mathcal{A}} + \lambda \boldsymbol{w}_{\alpha^*} = 0 \tag{H.117}$$

which, upon multiplying $1 - \theta$ on both sides, yields:

$$(1 - \theta)\nabla_{\boldsymbol{w}_{\alpha^*}, \mathcal{A}} + (1 - \theta)\lambda \boldsymbol{w}_{\alpha^*} = 0 \tag{H.118}$$

Then, when $\theta \in (0, 1)$, we have:

$$\boldsymbol{w}_{\alpha(\theta)} := \arg\min_{\boldsymbol{w} \in \mathcal{W}, ||\boldsymbol{w}||_2 \leq C} \theta \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1 - \theta)\mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r) + \frac{\lambda}{2}||\boldsymbol{w}||_2^2 \tag{H.119}$$

which yields the first order condition:

$$\theta \nabla_{\boldsymbol{w}_{\alpha(\theta)}, \mathcal{K}} + (1 - \theta)\nabla_{\boldsymbol{w}_{\alpha(\theta)}, \mathcal{A}} + \lambda \boldsymbol{w}_{\alpha(\theta)} = 0 \tag{H.120}$$

Subtracting Eq. (H.118) from Eq. (H.120) yields:

$$\theta \nabla_{\boldsymbol{w}_{\alpha(\theta)},\mathcal{K}} + (1-\theta)\nabla_{\boldsymbol{w}_{\alpha(\theta)},\mathcal{A}} + \lambda \boldsymbol{w}_{\alpha(\theta)} - (1-\theta)\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{A}} - (1-\theta)\lambda \boldsymbol{w}_{\alpha^*} = 0 \tag{H.121}$$

which simplifies to:

$$\theta \nabla_{\boldsymbol{w}_{\alpha(\theta)},\mathcal{K}} + (1-\theta)(\nabla_{\boldsymbol{w}_{\alpha(\theta)},\mathcal{A}} - \nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{A}}) + \lambda(\boldsymbol{w}_{\alpha(\theta)} - \boldsymbol{w}_{\alpha^*}) = -\theta\lambda \boldsymbol{w}_{\alpha^*} \tag{H.122}$$

The fundamental theorem of calculus then yields:

$$\int_0^1 \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*}+t(\boldsymbol{w}_{\alpha(\theta)}+\boldsymbol{w}_{\alpha^*}),\mathcal{A}}(\boldsymbol{w}_{\alpha(\theta} - \boldsymbol{w}_{\alpha^*})dt = \nabla_{\boldsymbol{w}_{\alpha(\theta)},\mathcal{A}} - \nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{A}} \tag{H.123}$$

and

$$\int_0^1 \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*}+t(\boldsymbol{w}_{\alpha(\theta)}+\boldsymbol{w}_{\alpha^*}),\mathcal{K}}(\boldsymbol{w}_{\alpha(\theta)} - \boldsymbol{w}_{\alpha^*})dt = \nabla_{\boldsymbol{w}_{\alpha(\theta)},\mathcal{K}} - \nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} \tag{H.124}$$

We thus denote:

$$\bar{\boldsymbol{H}}_\mathcal{K} := \int_0^1 \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*}+t(\boldsymbol{w}_{\alpha(\theta)}+\boldsymbol{w}_{\alpha^*}),\mathcal{K}}dt \tag{H.125}$$

$$\bar{\boldsymbol{H}}_\mathcal{A} := \int_0^1 \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*}+t(\boldsymbol{w}_{\alpha(\theta)}+\boldsymbol{w}_{\alpha^*}),\mathcal{A}}dt \tag{H.126}$$

$$\Delta \boldsymbol{w} := \boldsymbol{w}_{\alpha(\theta)} - \boldsymbol{w}_{\alpha^*} \tag{H.127}$$

Incorporating Eq. (H.123) and Eq. (H.124) into Eq. (H.122) then yields:

$$
\begin{aligned}
& \theta(\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} + \bar{\boldsymbol{H}}_\mathcal{K}\Delta\boldsymbol{w}) + (1-\theta)(\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{A}} + \bar{\boldsymbol{H}}_\mathcal{A}\Delta\boldsymbol{w} - \nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{A}}) + \lambda\Delta\boldsymbol{w} = -\theta\lambda\boldsymbol{w}_{\alpha^*} \\
\Longleftrightarrow \quad & \theta\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} + \theta\bar{\boldsymbol{H}}_\mathcal{K}\Delta\boldsymbol{w} + (1-\theta)\bar{\boldsymbol{H}}_\mathcal{A}\Delta\boldsymbol{w} + \lambda\Delta\boldsymbol{w} + \theta\lambda\boldsymbol{w}_{\alpha^*} = 0 \\
\Longleftrightarrow \quad & \big(\theta\bar{\boldsymbol{H}}_\mathcal{K} + (1-\theta)\bar{\boldsymbol{H}}_\mathcal{A} + \lambda\boldsymbol{I}\big)\Delta\boldsymbol{w} = -\theta\big(\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} + \lambda\boldsymbol{w}_{\alpha^*}\big) \\
\Longleftrightarrow \quad & \big(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I} + \theta(\bar{\boldsymbol{H}}_\mathcal{K} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}) + (1-\theta)(\bar{\boldsymbol{H}}_\mathcal{A} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{A}})\big)\Delta\boldsymbol{w} \\
& = -\theta\big(\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} + \lambda\boldsymbol{w}_{\alpha^*}\big) \\
\Longleftrightarrow \quad & \big(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I}\big)\Delta\boldsymbol{w} = -\big(\theta(\bar{\boldsymbol{H}}_\mathcal{K} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}) + (1-\theta)(\bar{\boldsymbol{H}}_\mathcal{A} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{A}})\big)\Delta\boldsymbol{w} \\
& \quad - \theta\big(\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} + \lambda\boldsymbol{w}_{\alpha^*}\big).
\end{aligned}
\tag{H.128}
$$

Then, note that:

$$||\bar{\boldsymbol{H}}_\mathcal{K} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 = ||\int_0^1 \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*}+t\Delta\boldsymbol{w},\mathcal{K}}dt - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 \tag{H.129}$$

$$\leq \int_0^1 ||\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*}+t\Delta\boldsymbol{w},\mathcal{K}} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 \tag{H.130}$$

$$\leq \frac{F_\mathcal{K}}{2}||\Delta\boldsymbol{w}||_2 \tag{H.131}$$

by the same token as in Prop. F.3.

Similarly:

$$||\bar{\boldsymbol{H}}_{\mathcal{A}} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{A}}||_2 \leq \frac{F_{\mathcal{A}}}{2}||\Delta\boldsymbol{w}||_2 \tag{H.132}$$

Also:

$$||\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 = ||\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}} - \nabla_{\mathcal{K}(\mathcal{D}_p),\mathcal{K}}||_2 \tag{H.133}$$
$$\leq P_{\mathcal{K}}||\boldsymbol{w}_{\alpha^*} - \boldsymbol{w}_{\mathcal{K}(\mathcal{D}_p)}||_2 \tag{H.134}$$
$$\leq 2P_{\mathcal{K}}C \tag{H.135}$$

by definition of the uniform learner $\mathcal{K}$, Lemma H.1, and the triangle inequality.

Additionally:

$$||\lambda\boldsymbol{w}_{\alpha^*}||_2 \leq \lambda C \tag{H.136}$$

By the triangle inequality, incorporating Eq. (H.131), Eq. (H.132), Eq. (H.135), and Eq. (H.136) into Eq. (H.128), we have that:

$$||(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})\Delta\boldsymbol{w}||_2 \leq (\theta||\bar{\boldsymbol{H}}_{\mathcal{K}} - \boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 + (1-\theta)||\bar{\boldsymbol{H}}_{\mathcal{A}} - \tag{H.137}$$
$$\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{A}}||_2)||\Delta\boldsymbol{w}||_2 + \theta||\nabla_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 + \theta||\lambda\boldsymbol{w}_{\alpha^*}||_2 \tag{H.138}$$
$$\leq (\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})||\Delta\boldsymbol{w}||_2^2 + \theta C(2P_{\mathcal{K}} + \lambda) \tag{H.139}$$

Note that we have, where $\sigma_{\min}(\cdot)$ denotes the minimum singular value:

$$||(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})\Delta\boldsymbol{w}||_2 \geq \sigma_{\min}(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})||\Delta\boldsymbol{w}||_2 \text{ by property of op. norm} \tag{H.140}$$
$$= \lambda_{\min}(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})||\Delta\boldsymbol{w}||_2 \text{ by Lemma H.3} \tag{H.141}$$

Furthermore, by Lemma H.1, we have that $||\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K}}||_2 \leq P_{\mathcal{K}}$ and $||\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{A}}||_2 \leq P_{\mathcal{A}}$, which yields:

$$||\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}}||_2 \leq \theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}} \tag{H.142}$$

which by Lemma H.3 yields:

$$\lambda_{\min}(\boldsymbol{H}_{\alpha^*,\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I}) \in [\lambda - \theta P_{\mathcal{K}} - (1-\theta)P_{\mathcal{A}}, \mu + \theta P_{\mathcal{K}} + (1-\theta)P_{\mathcal{A}}] \tag{H.143}$$

With Eq. (H.141), this yields that:

$$||(\boldsymbol{H}_{\boldsymbol{w}_{\alpha^*},\mathcal{K},\mathcal{A}} + \lambda\boldsymbol{I})\Delta\boldsymbol{w}||_2 \geq (\lambda - \theta P_{\mathcal{K}} - (1-\theta)P_{\mathcal{A}})||\Delta\boldsymbol{w}||_2 \tag{H.144}$$

Incorporating this into Eq. (H.139) yields:

$$(\lambda - \theta P_{\mathcal{K}} - (1-\theta)P_{\mathcal{A}})||\Delta\boldsymbol{w}||_2 \leq (\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})||\Delta\boldsymbol{w}||_2^2 + \theta C(2P_{\mathcal{K}} + \lambda) \tag{H.145}$$

Simplifying yields the quadratic inequality:

$$(\theta F_{\mathcal{K}} + (1-\theta)F_{\mathcal{A}})||\Delta\boldsymbol{w}||_2^2 - (\lambda - \theta P_{\mathcal{K}} - (1-\theta)P_{\mathcal{A}})||\Delta\boldsymbol{w}||_2 + \theta C(2P_{\mathcal{K}} + \lambda) \geq 0 \tag{H.146}$$

This then yields that:

$$||\Delta\boldsymbol{w}||_2 \leq \frac{\lambda - P - \sqrt{(\lambda - P)^2 - 4\theta CF(2P_{\mathcal{K}} + \lambda)}}{2F} \tag{H.147}$$

This is only valid when:

$$(\lambda - P)^2 - 4\theta CF(2P_{\mathcal{K}} + \lambda) \geq 0 \tag{H.148}$$

$$\iff \lambda^2 - 2L\lambda + P^2 - 4\theta C2P_{\mathcal{K}}F - 4\theta C\lambda F \geq 0 \tag{H.149}$$

$$\iff \lambda^2 - 2P\lambda - 4\theta CF\lambda + P^2 - 8\theta CP_{\mathcal{K}}F \geq 0 \tag{H.150}$$

$$\iff \lambda^2 - (2P + 4\theta CF)\lambda + (P^2 - 8\theta CP_{\mathcal{K}}F) \geq 0 \tag{H.151}$$

$$\impliedby \lambda \geq P + 2\theta CF + 2\sqrt{\theta CF(P + 2\theta CF + 8P_{\mathcal{K}}} \tag{H.152}$$

which holds by assumption; thus, Eq. (H.147) is valid. Note that all components of $2\theta CF + 2\sqrt{\theta CF(P + 2\theta CF + 8P_{\mathcal{K}}}$ are nonnegative, rendering this valid. Incorporating Eq. (H.147) into Lemma H.2 yields the final bound as desired. $\qquad\square$

Then, Theorem 5.5 follows as a corollary of Theorem H.13:

Proof. Note that we take care to ensure the bound holds for any choice of $\theta \in [0,1]$. Hence, fix $\theta \in [0,1]$.

Let

$$a := \lambda - P > 0, \qquad \varepsilon := 4\theta CF(2P_K + \lambda), \qquad \Delta := \frac{a - \sqrt{a^2 - \varepsilon}}{2F}. \tag{H.153}$$

Theorem H.13 gives the inequality:

$$|\alpha^* - \alpha(\theta)| \leq \frac{P_K}{2}\Delta^2 + \lambda C\Delta. \tag{H.154}$$

By the condition on λ in the theorem, the square root is real for every $\theta \in [0,1]$ (i.e. $a^2 - \varepsilon \geq 0$). This yields:

$$a - \sqrt{a^2 - \varepsilon} = \frac{\varepsilon}{a + \sqrt{a^2 - \varepsilon}}. \tag{H.155}$$

Then, since $a + \sqrt{a^2 - \varepsilon} \geq a > 0$, (H.155) implies:

$$a - \sqrt{a^2 - \varepsilon} \leq \frac{\varepsilon}{a}. \tag{H.156}$$

Dividing (H.156) by $2F$ yields:

$$\Delta \leq \frac{\varepsilon}{2aF}. \tag{H.157}$$

Then, substituting $\varepsilon = 4\theta CF(2P_K + \lambda)$ from (H.153) yields:

$$\Delta \leq \frac{4\theta CF(2P_K + \lambda)}{2aF} = \frac{2\theta C(2P_K + \lambda)}{a}. \tag{H.158}$$

We now bound the two terms on the right-hand side of (H.154). Using (H.158), we have that:

$$\lambda C \, \Delta \leq \lambda C \cdot \frac{2\theta C(2P_K + \lambda)}{a} = \frac{2\lambda(2P_K + \lambda)}{a} \, C^2 \theta = \mathcal{O}(\lambda C^2 \theta), \tag{H.159}$$

$$\frac{P_K}{2}\Delta^2 \leq \frac{P_K}{2}\left(\frac{2\theta C(2P_K + \lambda)}{a}\right)^2 = \frac{2P_K(2P_K + \lambda)^2}{a^2} \, C^2 \theta^2 = \mathcal{O}(C^2 \theta^2). \tag{H.160}$$

Combining (H.154), (H.159) and (H.160) and absorbing constants (which are independent of $\theta \in [0,1]$) yields:

$$|\alpha^* - \alpha(\theta)| = \mathcal{O}\big(\lambda C^2 \theta + C^2 \theta^2\big), \qquad \text{for any } \theta \in [0,1]. \tag{H.161}$$

as desired. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

I Experimental Details

I.1 Dataset Details

MNIST: The MNIST dataset contains 70k 28x28 greyscale images of hand-drawn digits in 10 classes (Deng, 2012). We conduct our experiments with 49k training images and 21k test images. The classes are mutually exclusive.

Kuzushiji-MNIST: The Kuzushiji-MNIST (KMNIST) dataset contains 70k 28x28 greyscale images of Japanese kanji in 10 classes (Clanuwat et al., 2018). We conduct our experiments with 49k training images and 21k test images. The classes are mutually exclusive.

CIFAR10: The CIFAR10 dataset consists of 60k 32x32 color images in 10 classes. The classes are mutually exclusive and include airplanes, automobiles, birds, cats, deer, dogs, frogs, horses, ships, and trucks (Krizhevsky et al., 2009).

CIFAR-100: The CIFAR-100 dataset is similar to CIFAR-10 but contains 100 classes, each with 600 images, making a total of 60k 32x32 color images. The 100 classes are grouped into 20 superclasses, and each image comes with a "fine" label (the class to which it belongs) and a "coarse" label (the superclass to which it belongs) (Krizhevsky et al., 2009).

SVHN: The Street View House Numbers (SVHN) dataset (Netzer et al., 2011) contains images of double-digit numbers on house walls as colored 32x32 images. We load SVHN with 10 classes. There are 73k training images, 26k testing images.

TinyImageNet: The TinyImageNet dataset (Le & Yang, 2015) contains 100000 64x64 images.

ImageNet: The ImageNet-1k dataset (Deng et al., 2009; Russakovsky et al., 2015) is a contains images of 1000 classes and approximately 1.2 million training images. The ImageNet-21k dataset, which our ViT models are pretrained with (Google, 2023), contains 21000 classes with approximately 14 million training images. Both datasets are 224x224.

I.2 Model Details

LogReg: A logistic regression model that has a single linear layer between inputs and outputs, followed by a softmax output function.

MLP: A two-layer ReLU feedforward neural network.

ResNet8: A [1,1,1,0] residual network, , with standard convolutional blocks, as described in (He et al., 2016).

ResNet18: A [2,2,2,2] residual network, with standard convolutional blocks, as described in (He et al., 2016).

ResNet50: A [3, 4, 6, 3] residual network, with bottleneck convolutional blocks, as described in (He et al., 2016).

ViT_S_16: A vision transformer with ≈ 20 million parameters, as detailed in (Dosovitskiy et al., 2021).

ViT_B_16: A vision transformer with ≈ 80 million parameters, as detailed in (Dosovitskiy et al., 2021).

I.3   Baseline Details

We implement several baselines and provide the rationale for their use below:

Pretrained: This is simply the pretrained model corresponding to whichever model and benchmark is specified. The rationale for using this is to demonstrate that we alter uniformity significantly from before without tarnishing accuracy for either the retain or test sets.

Retrained: This is a model retrained over the retain set, performing exact unlearning. The rationale for using this is to demonstrate that our methods mimic unlearning in how we preserve accuracy, but induce uniformity in a way that unlearning does not.

Synthetic: This method proceeds as follows: for each instance in the protected set, sample k instances from the ε-ball, with respect to the ℓ_2 norm, around that instance. Then, assign these k instances random labels from the label space, choosing labels uniformly at random. Do this for all protected set instances, yielding $|\mathcal{D}_p|k$ new instances. Then, append this to the retain set and retrain over this augmented dataset. This provides strong accuracy for simple baselines like MNIST while also inducing uniformity, providing an alternative, simple algorithm to compare our method against.

Label Differential Privacy: Specifically, we use the multi-stage training method of Ghazi et al. (2021) to obtain a model which is differentially private with respect to the labels–that is, an adversary cannot be sure whether the label they obtain is the true label. This is related to our work, albeit addresses a different threat model, as described in Appx. B. Still, we believe it is important to demonstrate that our method achieves privacy while not sacrificing utility to the extent that label differential privacy does, since it is a well-known method in the privacy literature that addresses a similar problem. For our experiments, we use the official repository with the hyperparameters reported in the paper: `https://github.com/google-research/label-dp`. Please note that we were unable to reproduce the test accuracy for CIFAR100. We contacted the authors and will update the manuscript after they respond. Regardless, LabelDP by definition will induce uniformity on the entire training dataset, at a higher cost to accuracy than Alg. 1, as observed for the other datasets (which we were able to reproduce accurately) in Appx. J.4.

I.4   Hyperparameter Details

Please note that, throughout, we do not do extensive hyperparameter optimization, which may lead to improved performance.

We use a standard train-test split of 70-30 throughout. Pretraining and synthetic training have the same hyperparameters as pretraining, unless mentioned otherwise. We use ADAM, with standard PyTorch hyperparameters aside from learning rate and weight decay, throughout. A batch size of 128 is used for pretraining and also for the retain set in Alg. 1 throughout. To ensure that gradients were well-aligned between the two conflicting objectives in Alg. 1, we implemented and ran gradient surgery (Yu et al., 2020) before each gradient step. Roughly, gradient surgery projects the gradient of one objective onto the normal plane of the other objective's gradient only when the two gradients conflict, i.e. have negative dot product. The projected gradient is chosen uniformly at random between the two. This effectively removes the component of the gradient that would increase the other objective's loss. For all Alg. 1 experiments and Alg. 2 experiments, we use a protected set size of 100 with a batch size (when loading the protected set into the finetuning in Alg. 1) of 10. We perform Alg. 1 for 100 epochs and finetune Alg. 2 for 50 epochs before running the certified Newton step. We generally use the forward KL divergence between model softmax outputs and the uniform distribution for $\mathcal{L}_{\mathcal{K}}$ and the cross entropy between model predictions and ground truth labels for $\mathcal{L}_{\mathcal{A}}$. For LogReg, we instead use the square loss between the uniform softmax probabilities and the model softmax outputs, since the forward KL is not necessarily convex in $\boldsymbol{w}$ in this case, while the square loss is; this allows us to use Alg. 2 with small λ. For our synthetic baseline, we use $\varepsilon = 8/255$ throughout, where ε is the size of

the ε-ball where we sample instances to assign random labels for retraining. For the LabelDP baseline, we use the multi-stage training algorithm of Ghazi et al. (2021) throughout.

Early stopping is implemented by saving the model which first meets the early stopping conditions, and continuing to see if any model performs better in terms of confidence distance while still meeting the early stopping conditions specified below.

Compute: We use two RTX 6000 Ada Generation NVIDIA GPUs throughout. The most resource intensive experiments are the LabelDP experiments, which take up most of the memory on both GPUs. Besides those, the other experiments take up at most a fourth of the compute resources available on one GPU. No experiments ran required more compute than these two GPUs provide.

MNIST, LogReg Pretraining: Epochs: 25. Learning rate: 0.01.

MNIST, MLP Pretraining: Epochs: 5. Learning rate: 0.01.

MNIST, ResNet18 Pretraining: Epochs: 2. Learning rate: 0.001.

MNIST, LogReg Alg. 1: Learning rate: 0.01

MNIST, MLP Alg. 1: Learning rate: 0.01

MNIST, SVM Alg. 1: Epochs: 100. Learning rate: 0.1. Optimized with SGD.

MNIST, ResNet18 Alg. 1: Learning rate: 0.001. Early stopping criterion of a confidence distance < 0.32 and a retain accuracy of $> 90\%$.

MNIST, LogReg Alg. 2: $M = 1$. $C = 10$, pretrained with PGD with the same hyperparameters as the standard pretraining. Since the losses are convex in $\boldsymbol{w}$, $\lambda_{\min} = 0$. $\lambda = 0.0001$. Following Zhang et al. (2024), we use the variance σ^2 as a hyperparameter, corresponding to a broad range of choices of ε and δ. We choose $\sigma = 0.001$. This results in large ε when choosing small δ, as typical in differential privacy (Dwork et al., 2014) and certified unlearning (Qiao et al., 2025). However, we still observe good induced uniformity. We also use a learning rate of $\alpha = 0.5$ in the Newton step, which, following Appx. F and bounding the gradient with αF instead of F, results in roughly twice as large of an ε.

MNIST, LogReg Alg. 3: $\lambda = 1.5, J = 4.0, n = 1000, b = 100$, and otherwise same hyperparameters as Alg. 2.

MNIST, MLP Alg. 3: $\lambda = 1.4, J = 2.5, n = 1000, b = 100$, and a learning rate of $\alpha = 0.35$.

MNIST, SVM Alg. 3: $\lambda = 0.5, J = 2.5, n = 5000, b = 200$, and a learning rate of $\alpha = 0.05$. We also used a twice-differentiable surrogate for the hinge loss to replace the standard hinge loss, which was used as the pretraining loss $\mathcal{L}_{\mathcal{A}}$.

MNIST, LogReg Synthetic Baseline: Sampled k instances for each protected set instance: 5.

MNIST, ResNet18 Synthetic Baseline: Sampled k instances for each protected set instance: 500.

MNIST, LogReg LabelDP Baseline: Epochs: 200. Batch size: 256. Random flip, random left-right flip, and random cutout (8). SGD with learning rate 0.4 (piecewise constant learning rate scheduler, as described by Ghazi et al. (2021)) with momentum 0.9. $\varepsilon = 2.0$. Mixup for stage 1: 16. Mixup for stage 2: 8. Data split evenly between the two stages. These hyperparameters are chosen to match those in the best results of Ghazi et al. (2021). See Ghazi et al. (2021) for more details on these hyperparameters.

MNIST, ResNet18 LabelDP Baseline: Same as the MNIST LogReg LabelDP hyperparameters, except with a weight decay of 0.0001 throughout.

KMNIST, LogReg Pretraining: Epochs: 100. Learning rate: 0.01.

KMNIST, MLP Pretraining: Epochs: 100. Learning rate: 0.001.

KMNIST, ResNet18 Pretraining: Epochs: 12 Learning rate: 0.002.

KMNIST, LogReg Alg. 1: Same as pretraining.

KMNIST, MLP Alg. 1:   Learning rate: 0.01.

KMNIST, ResNet18 Alg. 1:   Learning rate: 0.002. Early stopping criterion of a confidence distance < 0.32 and a retain accuracy of $> 99\%$.

KMNIST, ResNet18 Synthetic Baseline: Sampled k instances for each protected set instance: 500.

KMNIST, ResNet18 LabelDP Baseline: Same as the MNIST ResNet18 LabelDP hyperparameters.

SVHN, ResNet50 Pretraining: Epochs: 150. Learning rate: 0.001. Weight decay: 0.00005.

SVHN, ResNet50 Alg. 1:   Same as pretraining.

SVHN, ResNet50 Synthetic Baseline: Sampled k instances for each protected set instance: 500.

SVHN, ResNet50 LabelDP Baseline: Same as MNIST ResNet18 LabelDP hyperparameters.

CIFAR10, ResNet18 Pretraining: Epochs: 200 with SGD with a momentum of 0.9. Learning rate: 0.1. Weight decay: 0.0005.

CIFAR10, ResNet18 Alg. 1:   Same as pretraining. Early stopping criterion of a confidence distance < 0.42 and a retain accuracy of $> 85\%$.

CIFAR10, ResNet50 Pretraining: Same as CIFAR10 ResNet18.

CIFAR10, ResNet50 Alg. 1:   Same as pretraining. Early stopping criterion of a confidence distance < 0.42 and a retain accuracy of $> 85\%$.

CIFAR10, ResNet8 Pretraining: Same as CIFAR10 ResNet18.

CIFAR10, ResNet8 Alg. 1:   Same as CIFAR10 ResNet18.

CIFAR10, ResNet18 Synthetic Baseline: Sampled k instances for each protected set instance: 5000.

CIFAR10, ResNet50 Synthetic Baseline: Sampled k instances for each protected set instance: 5000.

CIFAR10, ResNet18 LabelDP Baseline: Same as MNIST ResNet18 LabelDP hyperparameters.

CIFAR10, ResNet50 LabelDP Baseline: Same as CIFAR10 ResNet18 LabelDP.

CIFAR10, ViT_S_16 Finetuning: 8 epochs. Learning rate 0.0001 with AdamW.

CIFAR10, ViT_B_16 Finetuning: 10 epochs. Learning rate 0.0001 with AdamW.

CIFAR100, ResNet50 Pretraining: Same as CIFAR10 ResNet18.

CIFAR100, ResNet50 Alg. 1:   Same as pretraining. Early stopping criterion of a confidence distance < 0.42 and a retain accuracy of $> 87\%$.

CIFAR100, ResNet8 Pretraining: Same as CIFAR100 ResNet50.

CIFAR100, ResNet8 Alg. 1:   Same as CIFAR100 ResNet50.

CIFAR100, ResNet50 Synthetic Baseline: Sampled instances: 5000.

CIFAR100, ResNet50 LabelDP Baseline: Same as CIFAR10 ResNet18 LabelDP, except with a batch size of 256 and a peak learning rate of 0.2. Please note that our results differ from the results reported in the original paper of Ghazi et al. (2021); however, we verified our results through several runs and used the official paper repository at `https://github.com/google-research/label-dp` with the hyperparameters reported in the paper.

CIFAR100, ViT_S_16 Finetuning: 30 epochs. Learning rate 0.002 with SGD with momentum 0.9. 500 warmup steps with cosine scheduler.

CIFAR100, ViT_S_16 Alg. 1:   100 epochs. Learning rate 0.001 with SGD.

CIFAR100, ViT_B_16 Finetuning: 45 epochs. Learning rate 0.002 with SGD with momentum 0.9. 500 warmup steps with cosine scheduler.

Figure J.5: Accuracy on test set for baselines as well as Alg. 1 and Alg. 2 with $\theta = 0.75$.

CIFAR100, ViT_B_16 Alg. 1: Same as CIFAR100 ViT_S_16.

TinyImageNet, ViT_S_16 Finetuning: 30 epochs. Learning rate 0.0001, momentum 0.9, and weight decay 0.01 with SGD.

TinyImageNet, ViT_S_16 Alg. 1: Same as CIFAR100 ViT_S_16.

TinyImageNet, ViT_B_16 Finetuning: 50 epochs. Learning rate 0.0001, momentum 0.9, and weight decay 0.01 with SGD.

TinyImageNet, ViT_B_16 Alg. 1: Same as CIFAR100 ViT_S_16.

ImageNet, ViT_B_16 Alg. 1: We load a pretrained model from Appx. J.3. We then finetune with Alg. 1 with SGD (0.9 momentum) and with learning rate 0.05. We finetune the Pareto learner for 22 epochs total, due to computational constraints; still, we obtain excellent performance as demonstrated in Tab. J.3.

Attacks: $\alpha = 0.0001$. PGD learning rate: 0.001. PGD steps: 50.

Test-Set Finetuning: Finetune pretrained model for 20 more epochs with the same hyperparameters as pretraining, then run Alg. 1 with the same hyperparameters as the original run of Alg. 1. For CIFAR10/CIFAR100, finetune for 100 epochs.

J   Additional Experiments

J.1   Test Set Accuracies for Main Paper Experiments

We provide a figure, similar to Fig. 2a, for the test set in Fig. J.5. We observe similar results as one does on the retain set, with test accuracies preserved by Alg. 1 and Alg. 2.

J.2   Tables for Main Paper Experiments

The tables for Fig. 2 are in Tab. J.1 and Tab. J.2. We include results for MNIST and KMNIST ResNet18 as well. We see that Alg. 1 induces uniformity without great damage to utility, while all other baselines–including the synthetic baseline–fail to do so without critically harming utility. Furthermore, we observe that ResNet50 performs better than ResNet18, providing more credibility to the claim in Sec. 6 that larger models tend to perform better when used in Alg. 1. Next, we observe that Alg. 1 can actually provide better retain and test accuracy than the pretrained model, as observed for ResNet18 over KMNIST; this is because we also

Table J.1: Results for Alg. 1, used in Fig. 2. We find that we are able to induce uniformity while only slightly decreasing retain and test accuracy. $\theta = 0.75$ throughout.

Dataset	Model	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
MNIST	ResNet18	Pretrain	98.0%	98.1%	0.877
		Retrain	97.3%	97.1%	0.876
		Synthetic	100.0%	99.1%	0.010
		LabelDP	98.8%	98.8%	0.593
		Alg. 1	99.6%	99.1%	0.070
KMNIST	ResNet18	Pretrain	98.2%	92.1%	0.880
		Retrain	98.4%	92.4%	0.884
		Synthetic	99.9%	96.7%	0.019
		LabelDP	98.9%	96.1%	0.530
		Alg. 1	99.1%	94.7%	0.257
SVHN	ResNet50	Pretrain	$100.0\% \pm 0.0\%$	$95.7\% \pm 0.1\%$	0.900 ± 0.000
		Retrain	$100.0\% \pm 0.0\%$	$95.7\% \pm 0.2\%$	0.885 ± 0.005
		Synthetic	$99.7\% \pm 0.1\%$	$93.7\% \pm 0.7\%$	0.873 ± 0.015
		LabelDP	$91.3\% \pm 0.5\%$	$92.2\% \pm 0.7\%$	0.280 ± 0.014
		Alg. 1	$\textcolor{red}{99.8\% \pm 0.2\%}$	$\textcolor{red}{95.1\% \pm 0.5\%}$	$\textcolor{red}{0.296 \pm 0.080}$
CIFAR10	ResNet18	Pretrain	100.0%	95.3%	0.898
		Retrain	100.0%	95.3%	0.891
		Synthetic	94.0%	89.7%	0.844
		LabelDP	85.8%	83.6%	0.359
		Alg. 1	89.6%	83.1%	0.377
	ResNet50	Pretrain	$100.0\% \pm 0.0\%$	$95.4\% \pm 0.4\%$	0.900 ± 0.000
		Retrain	$100.0\% \pm 0.0\%$	$95.3\% \pm 0.4\%$	0.881 ± 0.012
		Synthetic	$91.8\% \pm 1.0\%$	$87.9\% \pm 0.5\%$	0.803 ± 0.018
		LabelDP	$83.1\% \pm 0.5\%$	$81.0\% \pm 0.8\%$	0.3343 ± 0.012
		Alg. 1	$\textcolor{red}{87.9\% \pm 4.2\%}$	$\textcolor{red}{84.7\% \pm 3.3\%}$	$\textcolor{red}{0.359 \pm 0.007}$
CIFAR100	ResNet50	Pretrain	$100.0\% \pm 0.0\%$	$78.7\% \pm 1.2\%$	0.904 ± 0.004
		Retrain	$100.0\% \pm 0.0\%$	$79.0\% \pm 0.4\%$	0.766 ± 0.039
		Synthetic	$42.1\% \pm 2.3\%$	$40.6\% \pm 1.8\%$	0.365 ± 0.069
		LabelDP	$8.2\% \pm 1.0\%$	$7.8\% \pm 1.1\%$	0.115 ± 0.021
		Alg. 1	$\textcolor{red}{88.7\% \pm 13.4\%}$	$\textcolor{red}{66.9\% \pm 11.8\%}$	$\textcolor{red}{0.292 \pm 0.123}$

minimize the retain loss during finetuning. We similarly have to use early stopping for Alg. 1, as discussed in Sec. 6, since we use large models. Finally, we observe that LabelDP can induce uniformity, albeit at the cost of retain and test accuracy, but does so not only on the protected set but also the retain set; for a comparison of the confidence distances across the retain, test, and protected sets for LabelDP and our method, please see Tab. J.4. Additionally, for larger, more complex datasets like CIFAR100, LabelDP fails entirely. Please note that we do not perform extensive hyperparameter optimization during pretraining or retraining. Finally, note that our test set accuracy drop for a ResNet50 trained on CIFAR100 is unusually large; as demonstrated in Tab. J.3, training with larger models mitigate this drop.

We observe similar results for Alg. 2 in Tab. J.2. Additionally, we include t-statistic confidence intervals, with 4 degrees of freedom, across 5 seeds in J.1 for SVHN ResNet50, CIFAR10 ResNet50, and CIFAR100 ResNet50, as well as Tab. J.2. We find that we are able to systematically induce uncertainty on the protected set, while still keeping accuracy on retain and test sets.

J.3 Additional Experiments on TinyImageNet, ImageNet, & ViT

We provide experimental results for Alg. 1 for ViT trained on CIFAR100, TinyImageNet, and ImageNet in Tab. J.3, observing similar behavior–in fact significantly lower confidence distance with little retain or test

Table J.2: Results for Alg. 2 for logistic regression trained over MNIST, used in Fig. 2. $\theta = 0.75$ throughout.

Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
Pretrain	$92.7\% \pm 0.0\%$	$92.4\% \pm 0.1\%$	0.796 ± 0.025
Retrain	$92.7\% \pm 0.0\%$	$92.4\% \pm 0.1\%$	0.783 ± 0.015
Synthetic	$92.8\% \pm 0.4\%$	$91.1\% \pm 0.4\%$	0.726 ± 0.024
LabelDP	$56.7\% \pm 2.2\%$	$57.3\% \pm 2.4\%$	0.131 ± 0.016
Alg. 1	$89.2\% \pm 0.5\%$	$89.0\% \pm 0.4\%$	0.362 ± 0.079
Alg. 2	$82.2\% \pm 7.0\%$	$82.3\% \pm 7.1\%$	0.328 ± 0.085

Table J.3: Results for Alg. 1 for ViT trained on CIFAR100 and TinyImageNet.

Dataset	Model	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
CIFAR100	ViT_S_16	Pretrain	95.2%	90.1%	0.942
		Retrain	93.4%	89.1%	0.883
		Synthetic	86.4%	84.8%	0.682
		Alg. 1	91.6%	85.1%	0.036
	ViT_B_16	Pretrain	94.2%	91.2%	0.972
		Retrain	95.2%	91.0%	0.952
		Synthetic	17.7%	17.6%	0.189
		Alg. 1	91.6%	88.6%	0.074
TinyImageNet	ResNet50	Pretrain	91.2%	83.5%	0.812
		Retrain	91.6%	80.9%	0.924
		Synthetic	92.1%	80.6%	0.569
		Alg. 1	92.1%	81.6%	0.197
	ViT_S_16	Pretrain	86.7%	84.4%	0.698
		Retrain	87.2%	84.4%	0.742
		Synthetic	87.9%	86.3%	0.833
		Alg. 1	84.2%	81.4%	0.057
	ViT_B_16	Pretrain	95.0%	91.7%	0.822
		Retrain	96.8%	90.6%	0.826
		Synthetic	98.0%	84.4%	0.830
		Alg. 1	91.8%	88.3%	0.037
ImageNet	ViT_B_16	Pretrain	88.2%	81.1%	0.803
		Alg. 1	88.0%	80.7%	0.037

accuracy reduction–when compared to in Tab. J.1. Importantly, our results for a ViT trained on ImageNet yield nearly perfect uniformity with $< 1\%$ decreases in retain/test accuracy.

J.4 LabelDP and Alg. 1 Confidence Distances for Retain, Test, and Protected Sets

Here, we present Tab. J.4, which details the confidence distances for the retain, test, and protected sets of our method vs. LabelDP. Not only do we achieve better retain and test accuracy, but also we induce uniformity on *only* the protected set, while **LabelDP, undesirably, induces uniformity on the protected, retain, and test sets** altogether, functionally the same as adjusting the temperature. This does not suffice for our threat model, since we want to preserve confident predictions on the retain and test sets.

J.5 Pareto Frontier Main Paper Table

The results which correspond to Fig. 3b, Fig. 3a, and Fig. J.12 are included in Tab. J.5 and Tab. J.6.

Table J.4: A comparison of the confidence distances on the retain, test, and protected sets between Alg. 1 and LabelDP. In general, we induce uniformity on only the protected set, while maintaining confidently correct predictions on the retain and test sets, while LabelDP falls short. Note that this is only for one experiment run.

Dataset	Model	Method	Retain Conf. Dist. (Higher Better)	Test Conf. Dist. (Higher Better)	Protected Conf. Dist. (Lower Better)
MNIST	ResNet18	LabelDP	0.579	0.577	0.593
		Alg. 1	0.503	0.509	0.070
KMNIST	ResNet18	LabelDP	0.495	0.466	0.530
		Alg. 1	0.870	0.828	0.257
SVHN	ResNet50	LabelDP	0.288	0.285	0.282
		Alg. 1	0.888	0.855	0.280
CIFAR10	ResNet18	LabelDP	0.371	0.365	0.359
		Alg. 1	0.724	0.690	0.377
	ResNet50	LabelDP	0.366	0.361	0.334
		Alg. 1	0.725	0.701	0.270
CIFAR100	ResNet50	LabelDP	0.182	0.156	0.203
		Alg. 1	0.576	0.470	0.298

J.6 Optimization Dynamics

J.6.1 Empirical Results

Upon using Alg. 1, we observe that logistic regression fails to induce uniformity for more complex benchmarks than MNIST, e.g. KMNIST. Logistic regression has poor test accuracy; we thus conclude that a model must be large enough to generalize well in order to have uniformity induced over it without a large cost to retain accuracy. We discuss mathematical intuition for this in Appx. J.6.2.

However, when using Alg. 1 on larger models, we observe that we need early stopping. After achieving a good uniformity-utility tradeoff, large models e.g. ResNet50 on CIFAR10 will powerfully increase accuracy at the cost of uniformity. This is undesired behavior when compared to, for example, a ResNet8 trained on CIFAR10, where we initially increase uniformity at the cost of accuracy but slowly regain accuracy without critically damaging uniformity. We illustrate this in Fig. J.6. Altogether, our method works best for large models with early stopping during finetuning. We characterize this mathematically in Appx. J.6.2 as well.

We provide a plot characterizing how test accuracy for Alg. 1 and Alg. 2 applied on CIFAR10 and CIFAR100 for $\theta = 0.75$ changes over 100 epochs in Fig. J.7 for ResNet8 and ResNet50. We observe similar behavior to retain accuracy.

Results as used in Fig. J.6a, Fig. J.6b, and Fig. J.7 are included in Tab. J.7.

J.6.2 Intuition for Early Stopping

In what follows, we give mathematical justification for the behavior observed in Fig. J.6 and Tab. J.7.

Firstly, recall that random vectors are nearly orthogonal in high dimensions (Vershynin, 2018). In particular, for larger models, the gradients will conflict i.e. point in opposite directions more strongly, since their parameter space is very large. Second, every gradient step of our Alg. 1 is given by $\theta \nabla_{\boldsymbol{w}} \mathcal{L}_{\mathcal{K}}(\boldsymbol{w}, \mathcal{D}_p) + (1-\theta) \nabla_{\boldsymbol{w}} \mathcal{L}_{\mathcal{A}}(\boldsymbol{w}, \mathcal{D}_r)$.

In what follows, consider a large model e.g. CIFAR10 ResNet50.

When we begin finetuning the pretrained model with Alg. 1, θ is large, $\mathcal{L}_{\mathcal{A}}$ is small, and $\mathcal{L}_{\mathcal{K}}$ is large. Thus, $\nabla_{\boldsymbol{w}} \mathcal{L}_{\mathcal{K}}$ significantly dominates $\nabla_{\boldsymbol{w}} \mathcal{L}_{\mathcal{A}}$. For large models, since the gradients are nearly orthogonal, we will move very fast in the direction of the protected gradient.

Table J.5: Results for Alg. 1 and Alg. 2 as we explore the Pareto frontier over MNIST, single run for Fig. 3b, Fig. 3a, and Fig. J.12.

Model	θ	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
LogReg	0.000	93.7%	92.9%	0.817
	0.125	91.9%	91.2%	0.583
	0.250	92.9%	92.4%	0.178
	0.375	92.0%	92.3%	0.208
	0.500	92.4%	92.1%	0.342
	0.625	91.3%	91.1%	0.374
	0.750	91.6%	91.2%	0.263
	0.850	88.4%	87.5%	0.204
	0.950	89.2%	89.1%	0.169
Cert. LogReg	0.000	92.5%	92.2%	0.738
	0.125	92.2%	91.3%	0.573
	0.250	91.6%	91.5%	0.324
	0.375	91.4%	90.5%	0.327
	0.500	90.5%	90.6%	0.300
	0.625	90.6%	89.7%	0.451
	0.750	87.1%	87.2%	0.280
	0.850	89.3%	88.4%	0.206
	0.950	85.0%	85.7%	0.092
MLP	0.000	100.0%	98.1%	0.893
	0.125	98.3%	97.6%	0.399
	0.250	99.8%	97.4%	0.068
	0.375	97.4%	96.6%	0.247
	0.500	99.5%	97.1%	0.037
	0.625	96.4%	95.8%	0.166
	0.750	97.5%	95.7%	0.037
	0.850	92.5%	92.8%	0.096
	0.950	91.3%	90.0%	0.029
ResNet18	0.000	99.6%	99.4%	0.896
	0.125	99.3%	99.1%	0.892
	0.250	97.2%	97.5%	0.427
	0.375	99.6%	99.3%	0.586
	0.500	97.0%	97.0%	0.357
	0.625	97.0%	97.0%	0.357
	0.750	97.0%	97.0%	0.151
	0.850	95.7%	95.4%	0.234
	0.950	93.6%	94.0%	0.288

As finetuning continues, we reach a point where θ is large, $\mathcal{L}_{\mathcal{K}}$ is small, and $\mathcal{L}_{\mathcal{A}}$ is large. At this point, the $\nabla_{\boldsymbol{w}}\mathcal{L}_{\mathcal{A}}$ begins to dominate, albeit less significantly since θ is large. For large models, since the gradients are nearly orthogonal, we will move fast in the direction of the retain gradient. However, due to our choice of large θ, the retain gradient will be reduced in magnitude. Thus, we will move more slowly at this stage.

We must stop shortly after this, otherwise the protected loss will climb back up to a point where the model is no longer reasonably uniform.

Importantly, since smaller models (e.g. CIFAR10 ResNet8) have smaller parameter spaces, the gradients do not conflict as much. Thus, when we begin finetuning the model, while we do increase in retain loss initially, after the uniform loss is minimized we can minimize the retain loss freely. As such, we can move in a direction that minimizes both the protected and retain gradients, and do not need to stop early.

Table J.6: Results for Alg. 1 as we explore the Pareto frontier over CIFAR10 and CIFAR100 for ResNet50, single run for Fig. 3b, Fig. 3a, and Fig. J.12.

Dataset	θ	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
CIFAR10	0.000	100.0%	92.9%	0.817
	0.125	99.9%	94.1%	0.616
	0.250	99.9%	93.8%	0.586
	0.375	99.9%	94.6%	0.501
	0.500	91.9%	85.7%	0.395
	0.625	90.9%	86.9%	0.343
	0.750	94.7%	90.0%	0.270
	0.850	56.8%	56.0%	0.108
	0.950	10.7%	10.4%	0.168
CIFAR100	0.000	92.5%	92.2%	0.738
	0.125	99.9%	77.0%	0.353
	0.250	99.9%	77.5%	0.252
	0.375	99.9%	77.0%	0.296
	0.500	99.9%	77.1%	0.301
	0.625	90.9%	70.2%	0.353
	0.750	91.4%	65.5%	0.298
	0.85	21.2%	20.6%	0.138
	0.950	40.8%	30.3%	0.019

(a) Retain Accuracy vs. Epochs, $\theta = 0.75$

(b) Confidence Distance vs. Epochs, $\theta = 0.75$

Figure J.6: For CIFAR10 and CIFAR100 ResNet50, we observe a sharp drop in confidence distance followed by a sharp increase in Fig. J.6b, in line with the drops and increases for retain accuracy in Fig. J.6a. Test accuracy is similar. This highlights the need for early stopping when using Alg. 1 for large models, since otherwise one escapes from a good privacy-utility tradeoff. For smaller models, e.g. MNIST MLP, this issue does not persist–we obtain good uniformity after an initial drop in accuracy, but then increase accuracy and decrease confidence distance simultaneously.

However, as noted in the main paper, the model needs to be sufficiently large to achieve strong test accuracy, e.g. logistic regression trained on KMNIST does not work well. We hypothesize that this is because a model with more parameters has many more subspaces where two task losses can coexist in a way that provides a good tradeoff. We leave studying this and the above more formally to future work.

Table J.7: Studying the optimization dynamics of Alg. 1. Used in Fig. J.6.

Dataset	Model	Epoch	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
CIFAR10	ResNet8	0	98.9%	90.8%	0.883
		10	88.9%	80.4%	0.713
		20	90.8%	83.1%	0.539
		30	92.9%	85.3%	0.386
		40	95.4%	87.3%	0.295
		50	96.5%	88.9%	0.263
		60	97.3%	89.9%	0.261
		70	97.9%	90.0%	0.295
		80	98.2%	90.4%	0.270
		90	98.6%	90.5%	0.276
		100	98.9%	90.9%	0.245
CIFAR10	ResNet50	0	100.0%	95.2%	0.899
		10	72.5%	66.4%	0.593
		20	77.9%	75.0%	0.176
		30	88.8%	85.3%	0.131
		40	92.9%	88.7%	0.231
		50	96.1%	91.1%	0.394
		60	98.0%	92.5%	0.459
		70	98.9%	93.2%	0.579
		80	99.4%	93.7%	0.636
		90	99.6%	94.0%	0.665
		100	99.7%	94.0%	0.665
CIFAR100	ResNet8	0	96.4%	67.8%	0.515
		10	84.9%	59.2%	0.112
		20	89.5%	63.4%	0.063
		30	93.0%	67.3%	0.048
		40	94.6%	68.5%	0.046
		50	95.2%	69.0%	0.038
		60	95.6%	69.4%	0.036
		70	95.9%	69.5%	0.031
		80	96.0%	69.8%	0.030
		90	96.1%	70.0%	0.026
		100	96.2%	70.2%	0.026
CIFAR100	ResNet50	0	100.0%	78.3%	0.906
		10	64.9%	50.8%	0.170
		20	83.6%	65.6%	0.348
		30	93.1%	71.7%	0.419
		40	96.9%	74.4%	0.467
		50	98.7%	75.5%	0.505
		60	99.4%	76.4%	0.519
		70	99.7%	76.7%	0.530
		80	99.8%	76.5%	0.546
		90	99.8%	76.8%	0.555
		100	99.9%	77.3%	0.561

Figure J.7: Test Accuracy vs. Epochs, $\theta = 0.75$. This has similar behavior to Fig. J.6a.

Figure J.8: We corrupt an instance to increase the confidence of the final prediction. Noise is highlighted throughout for clarity.

J.7 Evaluating Test-Time Privacy Attacks

In what follows, we assume a test-time privacy (TTP) adversary with open-weight model access. We describe our attacks, specifically in Appx. A.

The results for the attacks with Alg. 4, Alg. 5, and Alg. 6 are in Tab. J.8 and Tab. J.9. We see that our method effectively defends against Alg. 4, Alg. 5, and Alg. 6 for various choices of γ in most scenarios. In our choices of γ, we follow the adversarial robustness literature, choosing γ sufficiently small such that the perturbation is invisible to the naked eye, but sufficiently large such that our attack is effective. However, in some cases, the attacks succeed despite the use of our method. Still, as demonstrated in Tab. J.10, we find that our algorithm renders the protected accuracy very low in these cases. As such, the adversary cannot be confidently correct–rather, in most cases, they can at best be confidently wrong. In particular, we have significantly better protection from attacks than in the pretrain, retrain, synthetic, or LabelDP cases. We provide visual intuition for our attacks in Fig. J.8.

J.8 Robustness of Alg. 1 Classifier on Neighboring Test Instances

To illustrate what happens for the finetuned classifier on neighboring test instances, we run an experiment evaluating accuracy and average confidence distance on test instances which are nearest neighbors to protected

---

Algorithm 4 Gaussian Noise Open-Weight Test Time Privacy Attack

---

Require: Protected set $\mathcal{D}_p$; pretrained model $\boldsymbol{w}^* = \mathcal{A}(\mathcal{D})$; adversarial γ

$\quad \mathcal{D}_p^{adv} = []$

$\quad$**for** $i = 1, ..., |\mathcal{D}_p|$ **do**

$\quad\quad \boldsymbol{x}_{adv} = \mathcal{D}_p^{(i)} + \beta,\ \beta \sim \mathcal{N}(0, \gamma\boldsymbol{I})$

$\quad\quad \mathcal{D}_p^{adv,(i)} = \boldsymbol{x}_{\text{adv}}$

$\quad$**end for**

$\quad\quad$**return** $\mathcal{D}_p^{\text{adv}}$

Algorithm 5 FGSM-Style Open-Weight Test Time Privacy Attack

Require: Protected set $\mathcal{D}_p$; pretrained model $\boldsymbol{w}^* = \mathcal{A}(\mathcal{D})$; adversarial γ; symmetry breaking α

$\quad \mathcal{D}_p^{adv} = []$

$\quad$**for** $i = 1, ..., |\mathcal{D}_p|$ **do**

$\quad\quad \boldsymbol{x}_0 = \mathcal{D}_p^{(i)} + \beta,\ \beta \sim U([-\alpha\gamma, \alpha\gamma])$

$\quad\quad \boldsymbol{x}_{\text{adv}} = \mathcal{D}_p^{(i)} + \gamma\text{sign}(\nabla_{\boldsymbol{x}}\rho(f_{\boldsymbol{w}^*}(\boldsymbol{x}))\big|_{\boldsymbol{x}=\boldsymbol{x}_0})$

$\quad\quad \mathcal{D}_p^{adv,(i)} = \boldsymbol{x}_{\text{adv}}$

$\quad$**end for**

$\quad\quad$**return** $\mathcal{D}_p^{\text{adv}}$

---

instances. Specifically, for each protected instance $\boldsymbol{x} \in \mathcal{D}_p$, we found the nearest neighbor of $\boldsymbol{x}$ in the test set. We evaluated this nearest neighbor in pixel space with respect to the ℓ_2 distance.

Results are in Tab. J.11. We notice that the classifier works as intended. That is, we obtain high accuracy as well as high confidence distance on the test set, including for nearby instances. While we do observe a drop for CIFAR100 ResNet50, we also observe that the confidence distance is still much higher than the 0.298 confidence distance on the protected set.

J.9   Ensuring Test-Time Privacy for Test Instances

In our paper, we focus on training data examples because this is the basis for scenarios addressed by the GDPR, HIPAA, etc. Providing the same guarantee for non-training (test) data is an equally important problem. However, the proposed method can be extended to cover this new case without loss of generality. One can just finetune on test instances highlighted to be corrupted and then run Alg. 1. In Tab. J.12, we find

---

Algorithm 6 PGD-Style Open-Weight Test Time Privacy Attack

Require: Protected set $\mathcal{D}_p$; pretrained model $\boldsymbol{w}^* = \mathcal{A}(\mathcal{D})$; adversarial γ; symmetry breaking α; step count N

$\quad \mathcal{D}_p^{adv} = []$

$\quad$**for** $i = 1, ..., |\mathcal{D}_p|$ **do**

$\quad\quad \boldsymbol{x}_{adv}^0 = \mathcal{D}_p^{(i)} + \beta,\ \beta \sim U([-\alpha\gamma, \alpha\gamma])$

$\quad\quad$**for** $j = 1, ..., N$ **do**

$\quad\quad\quad \boldsymbol{x}_{adv}^j = \boldsymbol{x}_{adv}^{j-1} + \beta\text{sign}(\nabla_{\boldsymbol{x}}\rho(f_{\boldsymbol{w}^*}(\boldsymbol{x}))\big|_{\boldsymbol{x}=\boldsymbol{x}_{\text{adv}}^{j-1}})$

$\quad\quad\quad \boldsymbol{x}_{adv}^j = \prod_{\mathcal{B}_\gamma(\mathcal{D}_p^{(i)})}(\boldsymbol{x}_{adv})$

$\quad\quad$**end for**

$\quad\quad \mathcal{D}_p^{adv,(i)} = \boldsymbol{x}_{\text{adv}}^N$

$\quad$**end for**

$\quad\quad$**return** $\mathcal{D}_p^{\text{adv}}$

Table J.8: Confidence distances over the Protected set after Alg. 4, Alg. 5, and Alg. 6 are applied to pretrained models and models finetuned with Alg. 1 for MNIST and KMNIST. Lower is better.

Dataset	Model	γ	Method	Attack	Prior Conf. Dist. (Lower Better)	Attack Conf. Dist. (Lower Better)
MNIST	MLP	$\frac{2}{255}$		Pretrain	Alg. 5	0.879	0.884
				Alg. 4	0.037	0.043
			Alg. 1	Alg. 5	0.037	0.054
				Alg. 6	0.037	0.185
		$\frac{5}{255}$		Pretrain	Alg. 5	0.879	0.888
				Alg. 4	0.037	0.064
			Alg. 1	Alg. 5	0.037	0.089
				Alg. 6	0.037	0.488
		$\frac{8}{255}$		Pretrain	Alg. 5	0.879	0.891
				Alg. 4	0.037	0.091
			Alg. 1	Alg. 5	0.037	0.125
				Alg. 6	0.037	0.632
	ResNet18	$\frac{2}{255}$		Pretrain	Alg. 5	0.895	0.896
				Alg. 4	0.070	0.075
			Alg. 1	Alg. 5	0.070	0.133
				Alg. 6	0.070	0.133
		$\frac{5}{255}$		Pretrain	Alg. 5	0.895	0.897
				Alg. 4	0.070	0.088
			Alg. 1	Alg. 5	0.070	0.164
				Alg. 6	0.070	0.350
		$\frac{8}{255}$		Pretrain	Alg. 5	0.895	0.898
				Alg. 4	0.070	0.116
			Alg. 1	Alg. 5	0.070	0.248
				Alg. 6	0.070	0.638
KMNIST	ResNet18	$\frac{2}{255}$		Pretrain	Alg. 5	0.858	0.865
				Alg. 4	0.257	0.258
			Alg. 1	Alg. 5	0.257	0.302
				Alg. 6	0.257	0.317
		$\frac{5}{255}$		Pretrain	Alg. 5	0.858	0.872
				Alg. 4	0.257	0.259
			Alg. 1	Alg. 5	0.257	0.370
				Alg. 6	0.257	0.420
		$\frac{8}{255}$		Pretrain	Alg. 5	0.858	0.878
				Alg. 4	0.257	0.259
			Alg. 1	Alg. 5	0.257	0.433
				Alg. 6	0.257	0.520

that finetuning with test instances yields similar performance to using our algorithm over just the training instances.

J.10 Ablation Study on Protected Set Size

Figures are provided in Fig. J.9 and Fig. J.10. In Tab. J.13, we provide experiments for ResNet50 trained on CIFAR10 and CIFAR100 for Alg. 1. In Tab. J.14, we provide experiments on MLP trained over MNIST for Alg. 1 and logistic regression trained over MNIST for Alg. 2.

Throughout our experiments, we use a protected set size of 100. We do so because for our use case, it is likely that a data controller would want to induce uniformity only for a small number of instances. We observe that as one increases the protected set size, it becomes harder to induce uniformity with the same

Table J.9: Confidence distances over the protected set after Alg. 4, Alg. 5, and Alg. 6 are applied to pretrained models and models finetuned with Alg. 1 for SVHN, CIFAR10, and CIFAR100. Lower is better.

Dataset	Model	γ	Method	Attack	Prior Conf. Dist. (Lower Better)	Attack Conf. Dist. (Lower Better)
SVHN	ResNet50	$\frac{1}{255}$	Pretrain	Alg. 5	0.972	0.886
			Alg. 1	Alg. 4	0.289	0.519
				Alg. 5	0.289	0.571
				Alg. 6	0.289	0.582
		$\frac{2}{255}$	Pretrain	Alg. 5	0.972	0.904
			Alg. 1	Alg. 4	0.289	0.519
				Alg. 5	0.289	0.613
				Alg. 6	0.289	0.640
CIFAR10	ResNet18	$\frac{1}{255}$	Pretrain	Alg. 5	0.898	0.898
			Alg. 1	Alg. 4	0.377	0.300
				Alg. 5	0.377	0.353
				Alg. 6	0.377	0.356
		$\frac{2}{255}$	Pretrain	Alg. 5	0.898	0.822
			Alg. 1	Alg. 4	0.377	0.301
				Alg. 5	0.377	0.397
				Alg. 6	0.377	0.408
	ResNet50	$\frac{1}{255}$	Pretrain	Alg. 5	0.900	0.806
			Alg. 1	Alg. 4	0.270	0.336
				Alg. 5	0.270	0.374
				Alg. 6	0.270	0.378
		$\frac{2}{255}$	Pretrain	Alg. 5	0.900	0.827
			Alg. 1	Alg. 4	0.270	0.336
				Alg. 5	0.270	0.407
				Alg. 6	0.270	0.425
CIFAR100	ResNet50	$\frac{1}{255}$	Pretrain	Alg. 5	0.906	0.587
			Alg. 1	Alg. 4	0.298	0.275
				Alg. 5	0.298	0.360
				Alg. 6	0.298	0.373
		$\frac{2}{255}$	Pretrain	Alg. 5	0.906	0.652
			Alg. 1	Alg. 4	0.298	0.276
				Alg. 5	0.298	0.421
				Alg. 6	0.298	0.468

hyperparameters. Still, for Alg. 1, we are able to obtain strong uniformity with good retain and test accuracy for significantly larger protected set sizes. Furthermore, we observe that Alg. 2 fails for sufficiently large protected set size; this is likely because Hessian matrix is significantly larger (in norm) for a larger protected set, resulting in a catastrophically large Newton step. Mitigating this phenomenon is left to future work, where Hessian-free techniques like those of Qiao et al. (2025) may be advantageous.

The results in figures Fig. J.9 and Fig. J.10, as well as Tab. J.14 and Tab. J.13, we keep the batch size used for the protected set in Alg. 1 to 10 instances. However, upon scaling the batch size proportionately with the protected set size (specifically dividing the protected set size by ten), we observe much better performance. These results are in Tab. J.15. Specifically, for a MLP trained on MNIST and a protected set size of 5000, we obtain much better uniformity than when we did not scale the batch size.

J.11 Evaluating Confidence Distance as a TTP Metric

The ℓ_2 metric $||f(\boldsymbol{x}) - \frac{\vec{1}}{K}||_2$ has similar utility to our presented metric. However, it is slightly less interpretable and may accidentally overpenalize uncertain outputs. For example, if one class has no probability but the

Table J.10: Accuracies over the protected set after Alg. 5 and Alg. 6 are applied to pretrained models and models finetuned with Alg. 1 for SVHN, CIFAR10, and CIFAR100. We see that Alg. 1 significantly lowers protected set accuracy.

Dataset	Model	γ	Method	Attack	Prior. Protected Acc. (Lower Better)	Atk. Protected Acc. (Lower Better)
MNIST	MLP	$\frac{5}{255}$	Pretrain	Alg. 6	97.0%	96.0%
			Alg. 1	Alg. 6	71.0%	43.0%
		$\frac{8}{255}$	Pretrain	Alg. 6	98.3%	96.0%
			Alg. 1	Alg. 6	71.0%	41.0%
	ResNet18	$\frac{8}{255}$	Pretrain	Alg. 6	98.9%	100.0%
			Alg. 1	Alg. 6	28.0%	54.0%
KMNIST	ResNet18	$\frac{8}{255}$	Pretrain	Alg. 6	96.0%	96.0%
			Alg. 1	Alg. 6	50.0%	55.0%
SVHN	ResNet50	$\frac{1}{255}$	Pretrain	Alg. 5	97.0%	66.0%
			Alg. 1	Alg. 4	40.0%	53.0%
				Alg. 5	40.0%	52.0%
				Alg. 6	40.0%	52.0%
		$\frac{2}{255}$	Pretrain	Alg. 5	97.0%	66.0%
			Alg. 1	Alg. 4	40.0%	52.0%
				Alg. 5	40.0%	53.0%
				Alg. 6	40.0%	53.0%
CIFAR10	ResNet18	$\frac{1}{255}$	Pretrain	Alg. 5	100.0%	61.0%
			Alg. 1	Alg. 4	60.0%	35.0%
				Alg. 5	60.0%	34.0%
				Alg. 6	60.0%	34.0%
		$\frac{2}{255}$	Pretrain	Alg. 5	100.0%	61.0%
			Alg. 1	Alg. 4	60.0%	35.0%
				Alg. 5	60.0%	33.0%
				Alg. 6	60.0%	32.0%
	ResNet50	$\frac{1}{255}$	Pretrain	Alg. 5	100.0%	56.0%
			Alg. 1	Alg. 4	55.0%	33.0%
				Alg. 5	55.0%	30.0%
				Alg. 6	55.0%	30.0%
		$\frac{2}{255}$	Pretrain	Alg. 5	100.0%	56.0%
			Alg. 1	Alg. 4	55.0%	30.0%
				Alg. 5	55.0%	30.0%
				Alg. 6	55.0%	32.0%
CIFAR100	ResNet50	$\frac{1}{255}$	Pretrain	Alg. 5	100.0%	32.0%
			Alg. 1	Alg. 4	48.0%	11.0%
				Alg. 5	48.0%	10.0%
				Alg. 6	48.0%	11.0%
		$\frac{2}{255}$	Pretrain	Alg. 5	100.0%	32.0%
			Alg. 1	Alg. 4	48.0%	11.0%
				Alg. 5	48.0%	10.0%
				Alg. 6	48.0%	11.0%

other 9 classes are uniform. Still, as demonstrated in Tab. J.16, we find that we minimize this metric as well for the same models and datasets. This holds similarly for other potential metrics e.g. the ℓ_1 metric.

Table J.11: Accuracies and confidence distances for test instances which are nearest neighbors (with respect to ℓ_2 distance) of protected set instances. We observe that models continue to confidently and correctly classify nearby test instances after finetuning with Alg. 1.

Dataset	Model	Method	Acc.	Conf. Dist. (Higher Better)
MNIST	MLP	Pretrain	100.0%	0.894
		Alg. 1	93.0%	0.754
	ResNet18	Pretrain	100.0%	0.896
		Alg. 1	100.0%	0.875
KMNIST	ResNet18	Pretrain	99.0%	0.871
		Alg. 1	99.0%	0.850
SVHN	ResNet50	Pretrain	95.0%	0.982
		Alg. 1	95.0%	0.939
CIFAR10	ResNet18	Pretrain	98.0%	0.876
		Alg. 1	85.0%	0.538
	ResNet50	Pretrain	96.0%	0.884
		Alg. 1	90.0%	0.700
CIFAR100	ResNet50	Pretrain	78.0%	0.778
		Alg. 1	66.0%	0.484

(a) Retain Accuracy vs. Epochs, $\theta = 0.75$ (b) Confidence Distance vs. Epochs, $\theta = 0.75$

Figure J.9: We observe that retain accuracy stays fairly stable as the protected size increases, in Fig. J.9a, except for Alg. 2 where it causes catastrophic failure due to the magnitude of the Newton step. Furthermore, we find that confidence distance slowly increases as the protected set size increases in Fig. J.9b.

J.12 Additional Baselines

J.12.1 Retraining with Randomly Sampled $\mathcal{D}_p$ Labels

Below, we present results for a baseline where we sample labels uniformly at random for the protected set and retrain with $\mathcal{D}_r \cup \mathcal{D}_u$. In Tab. J.17, we demonstrate that this fails to induce uncertainty, especially when compared to Alg. 1; this is expected, since well-trained models will maximize confidence on training instances by default.

J.12.2 Randomly Sampled Labels with GaussianUniform

In what follows, we present the *GaussianUniform* baseline, an alternative idea to our approach based on the notable work of Zhang et al. (2017), which demonstrates that a neural network can fully minimize its

Table J.12: Finetuning on test instances and running Alg. 1. Here, "pretrain" denotes the initially pretrained model (no additional test instances), while other rows correspond to finetuning on the specified number of test instances.

Dataset	Model	% Protected Size	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
MNIST	MLP	0	Pretrain	98.3%	97.0%	0.879
			Alg. 1	97.5%	95.7%	0.037
		1/5	Pretrain	99.1%	97.3%	0.892
			Alg. 1	96.4%	93.9%	0.147
		1/2	Pretrain	99.0%	97.1%	0.896
			Alg. 1	96.2%	93.4%	0.157
	ResNet18	0	Pretrain	99.2%	98.9%	0.879
			Alg. 1	99.6%	99.1%	0.070
		1/5	Pretrain	98.9%	98.5%	0.884
			Alg. 1	99.6%	98.9%	0.181
		1/2	Pretrain	99.4%	100.0%	0.897
			Alg. 1	99.5%	98.7%	0.243
KMNIST	ResNet18	0	Pretrain	98.2%	92.1%	0.880
			Alg. 1	99.1%	94.7%	0.257
		1/5	Pretrain	100.0%	97.5%	0.900
			Alg. 1	99.7%	96.5%	0.301
		1/2	Pretrain	100.0%	97.7%	0.895
			Alg. 1	99.5%	96.9%	0.233
SVHN	ResNet50	0	Pretrain	99.5%	94.4%	0.971
			Alg. 1	99.0%	94.4%	0.280
		1/5	Pretrain	99.7%	94.6%	0.974
			Alg. 1	99.2%	94.4%	0.184
		1/2	Pretrain	99.5%	94.5%	0.986
			Alg. 1	99.8%	95.1%	0.391
CIFAR10	ResNet18	0	Pretrain	100.0%	95.3%	0.898
			Alg. 1	89.6%	83.1%	0.377
		1/5	Pretrain	100.0%	98.4%	0.894
			Alg. 1	89.4%	85.6%	0.358
		1/2	Pretrain	100.0%	93.5%	0.887
			Alg. 1	87.9%	83.7%	0.494
	ResNet50	0	Pretrain	100.0%	95.2%	0.900
			Alg. 1	94.7%	90.0%	0.270
		1/5	Pretrain	99.9%	93.4%	0.896
			Alg. 1	91.6%	87.1%	0.392
		1/2	Pretrain	99.8%	93.6%	0.894
			Alg. 1	89.1%	86.7%	0.390
CIFAR100	ResNet50	0	Pretrain	100.0%	78.3%	0.902
			Alg. 1	91.4%	65.5%	0.298
		1/5	Pretrain	100.0%	79.0%	0.880
			Alg. 1	93.5%	69.1%	0.021
		1/2	Pretrain	99.9%	78.2%	0.860
			Alg. 1	97.9%	72.8%	0.221

loss over a training dataset where samples have labels sampled uniformly at random. The approach of GaussianUniform is as follows:

1. Begin with a training dataset $\mathcal{D} = \mathcal{D}_p \cup \mathcal{D}_r$.

Figure J.10: Test Accuracy vs. Protected Set Size, $\theta = 0.75$. This has similar behavior to Fig. J.9a.

Table J.13: Results on applying Alg. 1 on ResNet50 for various protected set sizes over CIFAR10 and CIFAR100. $\theta = 0.75$ throughout.

Dataset	Model	Protected Size	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
CIFAR10	ResNet50	100	94.7%	90.0%	0.270
		250	96.4%	90.3%	0.289
		500	95.1%	88.6%	0.190
		1000	97.0%	90.0%	0.144
		5000	95.8%	89.5%	0.176
CIFAR100	ResNet50	100	91.4%	65.5%	0.298
		250	99.7%	78.0%	0.370
		500	99.8%	76.1%	0.475
		1000	99.7%	74.8%	0.492
		5000	99.8%	74.1%	0.613

2. Perturb all samples in $\mathcal{D}$ to yield $\mathcal{D}' = \mathcal{D}'_p \cup \mathcal{D}'_r$. We use mean zero Gaussian noise with 0.1 variance, which adds a small amount of noise.

Table J.14: Results on applying Alg. 1 and Alg. 2 on various protected set sizes over MNIST. We observe that, while Alg. 1 still works well, confidence distance increases as protected set size does; please see Appx. J.10 for a discussion on Alg. 2. $\theta = 0.75$ throughout.

Method	Model	Protected Size	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
Alg. 1	MLP	100	97.5%	95.7%	0.037
		250	97.6%	95.1%	0.010
		500	95.8%	93.8%	0.121
		1000	94.8%	93.3%	0.162
		5000	96.6%	95.0%	0.420
Alg. 2	LogReg	100	92.1%	91.8%	0.258
		250	85.4%	85.2%	0.243
		500	5.7%	5.9%	0.824
		1000	4.4%	4.6%	0.891
		5000	2.4%	2.5%	0.899

Table J.15: Results on applying Alg. 1 on various protected set sizes over MNIST while scaling the batches. We observe much better performance for a protected set size of 5000 than in Tab. J.14, and comparable performance on other protected set sizes.

Method	Model	**Protected** Size	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
Alg. 1	MLP	100	97.5%	95.7%	0.037
		250	95.6%	96.6%	0.131
		500	93.8%	93.4%	0.187
		1000	92.3%	92.2%	0.192
		5000	95.3%	94.9%	0.227

Table J.16: ℓ_2 confidence distances for models finetuned with Alg. 1.

Dataset	Model	Conf. Dist. Type	Pretrained (Lower Better)	Alg. 1 (Lower Better)
MNIST	MLP	Paper	0.879	0.037
		ℓ_2	0.930	0.053
	ResNet18	Paper	0.895	0.070
		ℓ_2	0.944	0.070
KMNIST	ResNet18	Paper	0.880	0.257
		ℓ_2	0.911	0.302
SVHN	ResNet50	Paper	0.972	0.289
		ℓ_2	0.979	0.298
CIFAR10	ResNet18	Paper	0.898	0.377
		ℓ_2	0.947	0.435
	ResNet50	Paper	0.900	0.270
		ℓ_2	0.948	0.323
CIFAR100	ResNet50	Paper	0.902	0.298
		ℓ_2	0.911	0.311

3. Sample all labels in $\mathcal{D}$ uniformly at random to yield $\tilde{\mathcal{D}} = \tilde{\mathcal{D}}_f \cup \tilde{\mathcal{D}}_r$.

4. Train $\mathcal{A}(\mathcal{D}' \cup \tilde{\mathcal{D}})$.

In this scenario, inducing uncertainty may not be necessary, since the protected set would have very strong uniformity, with strong accuracy on the retain set available by slightly perturbing with Gaussian noise. We use the same hyperparameters as pretraining for the respective model and dataset tested, as reported in Appx. I. We find that this is not the case for ResNet50 trained on SVHN, CIFAR10, and CIFAR100. However, it achieves very poor test accuracy compared to Alg. 1 on both normal $\mathcal{D}_{\text{test}}$ and perturbed test $\mathcal{D}'_{\text{test}}$ datasets; thus, we prefer Alg. 1 to this approach, since it generalizes well and also does not require retraining. Results are in Tab. J.18.

J.13 Tightness of Bound in Theorem 5.5

To evaluate how tight our bound is, we run an experiment for MNIST logistic regression. We use the notation of Theorem 5.5 in Tab. J.19. We find that our constant bound is fairly tight as $\theta \to 1$; we leave using more advanced techniques to ensure better tightness to future work.

J.14 Proportions of Time Elapsed in Alg. 1

Results are reported in Tab. J.20.

Table J.17: Results for Alg. 1 compared to the "Relabel" baseline on ResNet18 and ResNet50 trained with MNIST and CIFAR10, respectively. We find that the "Relabel" baseline fails to induce uniformity, while our Alg. 1 does.

Dataset	Model	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
MNIST	ResNet18	Pretrain	98.0%	98.1%	0.877
		Relabel	100.0%	99.4%	0.893
		Alg. 1	99.6%	99.1%	0.070
CIFAR10	ResNet50	Pretrain	100.0%	95.2%	0.900
		Relabel	100.0%	95.3%	0.885
		Alg. 1	94.7%	90.0%	0.270
CIFAR100	ResNet50	Pretrain	100.0%	78.3%	0.902
		Relabel	99.95%	77.3%	0.730
		Alg. 1	91.4%	65.5%	0.298

Table J.18: Results for the *GaussianUniform* baseline described in Appx. J.12.2. This method results in significantly degraded accuracy on the test set compared to a model finetuned with Alg. 1, despite achieving high accuracy on the perturbed retain set (D'_r). Note that performance on $\tilde{\mathcal{D}}_r$ is similar to $\mathcal{D}'_r$.

Dataset	Model	Train Set Acc. (%)		Test Set Acc. (%)			Conf. Dist.
		$\mathcal{D}'_r$	$\mathcal{D}'_p$	$\mathcal{D}_{\text{test}}$	$\mathcal{D}'_{\text{test}}$	Alg. 1	$\mathcal{D}_p$
SVHN	ResNet50	82.4%	2.0%	6.2%	81.1%	94.4%	0.009
CIFAR10	ResNet50	100.0%	12.0%	11.1%	73.3%	90.0%	0.573
CIFAR100	ResNet50	99.9%	4.0%	3.4%	40.4%	65.5%	0.057

(a) Sample 1   (b) Sample 2   (c) Sample 3   (d) Sample 4   (e) Sample 5

Figure J.11: Comparison of Pretrain (red) and Alg. 1 (blue) softmax probabilities across five different CIFAR10 protected set samples.

J.15 Warmup Values for MNIST LogReg

Results are contained in Tab. J.21. We find that after applying the certified Newton step in Alg. 2, we obtain better retain and test accuracy, at small cost to uniformity. Thus, warming up is not the only component of achieving good results in Alg. 2.

J.16 Visualization of Softmax Outputs

We provide a comparison of pretrained f and Alg. 1 softmax probabilities across five different CIFAR10 protected set samples, demonstrating visually the effectiveness of Alg. 1 at inducing uniformity (and the relevance of our confidence distance metric) in Fig. J.11.

Table J.19: Comparison of bounds

$\alpha^* - \alpha(1)$	O(1) bound size
1.678	3.374

Table J.20: Time proportions of each step in Alg. 1.

Dataset	Model	Retain Grad.	Protected Grad.	Surgery	Reg. Grad.	Step
MNIST	MLP	0.967	0.033	0.000	0.000	0.000
	ResNet18	0.984	0.016	0.0010	0.000	0.000
KMNIST	ResNet18	0.989	0.011	0.000	0.000	0.000
SVHN	ResNet50	0.980	0.020	0.000	0.000	0.000
CIFAR10	ResNet18	0.989	0.011	0.000	0.000	0.000
	ResNet50	0.992	0.008	0.000	0.000	0.000
CIFAR100	ResNet50	0.993	0.006	0.000	0.000	0.000

Table J.21: Warmup values for Alg. 2 for logistic regression trained over MNIST, contrasted with the values after Alg. 2 is applied.

θ	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
0.0	Warmup	91.4%	91.6%	0.765
0.0	Alg. 2	92.5%	92.2%	0.738
0.25	Warmup	90.1%	90.4%	0.283
0.25	Alg. 2	91.6%	91.5%	0.324
0.50	Warmup	89.3%	89.7%	0.215
0.50	Alg. 2	90.5%	90.6%	0.300
0.75	Warmup	88.4%	88.6%	0.154
0.75	Alg. 2	87.1%	87.2%	0.280
0.95	Warmup	85.4%	86.0%	0.097
0.95	Alg. 2	85.0%	85.7%	0.092

J.17 Results for KMNIST LogReg and MLP

Results are contained in Tab. J.22. As mentioned in Sec. 6, one can see that a small model for a more complex benchmark, logistic regression on KMNIST, fails to induce uniformity, since the pretrained model is too small to generalize. However, on a bigger model i.e. an MLP trained over KMNIST, since it is large enough to generalize, one can induce uniformity over it. Thus, larger models which generalize well are preferred for our method, in line with the goals of ML.

J.18 Ablation Study on Synthetic Baseline Sample Size

Below, we study what we happen if we increase the number of samples sampled in the ε-ball in the synthetic baseline. For each protected set instance, we sample k instances from the ε-ball around the protected set, and then assign random labels to these instances, yielding an additional $|\mathcal{D}_p|k$ instances in the training data. We then retrain the model over the retain set along with these new $|\mathcal{D}_p|k$ instances. For a MLP trained over

Table J.22: Results for Alg. 1 applied to logistic regression and MLP trained over KMNIST, $\theta = 0.75$ for Alg. 1.

Model	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
LogReg	Pretrain	81.4%	66.4%	0.775
	Retrain	80.9%	65.3%	0.770
	Alg. 1	77.4%	63.4%	0.770
MLP	Pretrain	100%	88.4%	0.900
	Retrain	100%	88.5%	0.887
	Alg. 1	92.8%	80.3%	0.039

Table J.23: Results for the synthetic baseline applied with various sampled k on a MLP over MNIST and a ResNet18 over CIFAR10. We observe that increasing k yields better performance, but nevertheless even very large k (an additional 50k instances, with a protected set size of 100) fails to induce uniformity for CIFAR10.

Dataset	Model	Sampled k	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
MNIST	MLP	5	99.4%	97.1%	0.541
		25	99.0%	96.4%	0.183
		125	99.4%	96.8%	0.105
		250	98.6%	96.0%	0.066
		500	99.6%	96.3%	0.003
CIFAR10	ResNet18	5	99.0%	91.0%	0.683
		25	99.2%	91.3%	0.865
		125	98.8%	90.9%	0.856
		250	98.4%	90.7%	0.869
		500	98.3%	91.1%	0.852
		5000	94.0%	89.7%	0.844

Figure J.12: Test Accuracy vs. θ, MNIST. This has similar behavior to Fig. 3a.

MNIST, we observe better performance as we increase sample size. However, for a ResNet18 trained over CIFAR10, even if we have a very large sample size. This is presented in Tab. J.23. Thus, since Alg. 1 can induce uniformity as shown in Tab. J.1, without great cost to retain or test accuracy, it is better than the synthetic baseline.

J.19 Test Accuracy Plot for Pareto Frontier Experiments

We provide a plot characterizing test accuracy for Alg. 1 and Alg. 2 applied on MNIST for various choices of θ in Fig. J.12.

J.20 Results for Alg. 3

Due to the issues presented in Appx. E, we focus on Alg. 1 and Alg. 2 in the main paper and present results for Alg. 3 here. In Tab. J.24, we show that Alg. 3 can induce uniformity while maintaining accuracy for logistic regression, albeit with a stronger tradeoff than Alg. 2

Table J.24: Results for Alg. 3 on MNIST across different models. We compare against pretraining, Alg. 1, and Alg. 2 (where available). We observe that Alg. 3 achieves competitive uniformity on MLP, though with a trade-off in retain accuracy.

Dataset	Model	Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
MNIST	LogReg	Pretrain	92.1%	91.8%	0.807
		Alg. 1	87.8%	87.7%	0.180
		Alg. 2	87.1%	87.2%	0.280
		Alg. 3	77.9%	77.3%	0.267
	SVM	Pretrain	92.3%	92.1%	0.543
		Alg. 1	90.3%	90.1%	0.241
		Alg. 3	83.9%	83.9%	0.215
	MLP	Pretrain	98.3%	97.0%	0.886
		Alg. 1	97.5%	95.7%	0.037
		Alg. 3	88.8%	87.5%	0.057

(a) $\lambda = 1.0$ $J = 2.5$. Here, J is off by 0.5, and the estimator diverges.

(b) $\lambda = 0.5, J = 2.0$. Here, λ is off by 0.5, and the estimator diverges.

(c) $\lambda = 0.5, J = 2.5$. Here, the estimator stays roughly stable.

Figure J.13: We observe that small changes in hyperparameters result in divergence of the inverse Hessian estimator in Alg. 3.

J.21 Hyperparameter Sensitivity for Alg. 3

In what follows, we demonstrate the sensitivity of Alg. 3 to the choice of hyperparameters for logistic regression trained on MNIST (using the forward KL for uniformity). Here $b = 1$ and $n = 1000$. Below, the x-axis is the estimator recursive timesteps $t = 1, ..., n$ and the y axis is the difference between the exact Newton step and the estimated Newton step in Frobenius norm (before any Gaussian noise is added). Specifically, we find that if $\lambda \approx |\lambda_{\min}|$ and $J \approx |\lambda + \lambda_{\min}|$ does not hold, the estimator diverges very quickly. Specifically, as shown in Fig. J.13, if $\lambda = 1.0, J = 2.5$ (here, λ is off by 0.5) or $\lambda = 0.5, J = 2.0$ (here, J is off by 0.5), the estimator diverges before timestep 150, resulting in integer overflow. However, when $\lambda = 0.5, J = 2.5$, the estimator stays stable across samples and does not diverge. Still, the estimator does not converge but rather oscillates, demonstrating the poor sample complexity of the estimator as well–although it is asymptotically unbiased, it may take many samples to converge.

J.22 Test-Time Privacy in Text Classification

We provide a preliminary evaluation on text classification with the AG News (Del Corso et al., 2005), which contains four classes of news articles. We train on embedded representations of these, using a pretrained text embedder. We pretrain a MLP with ADAM, using a learning rate of 0.003 and a batch size of 256. In Tab. J.25, we find that Alg. 1 effectively induces uncertainty on this dataset as well, while still retaining high retain and test accuracies.

Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
Pretrain	93.3%	90.9%	0.645
Retrain	93.2%	90.9%	0.649
Alg. 1	91.9%	89.4%	0.221

Table J.25: We run Alg. 1 on a MLP trained on the AG News dataset, finding that we can induce uncertainty while keeping strong retain and test accuracies on text classification as well.

Method	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
Pretrain	98.0%	98.1%	0.877
Alg. 1	99.6%	99.1%	0.070
NegGrad+	93.8%	92.9%	0.862

Table J.26: We run the unlearning method NegGrad+ (Kurmanji et al., 2023) on a ResNet18 pretrained on MNIST with $\beta = 0.5$. We find that, while retain and test accuracy does not drop significantly, NegGrad+ does not induce protected set uncertainty while Alg. 1 does.

J.23 An Additional Unlearning Baseline

In Tab. J.26, we run an additional unlearning baseline NegGrad+ (Kurmanji et al., 2023). Similar to the retrain baseline in Tab. J.1, we find that NegGrad+ cannot induce uncertainty on the protected set, while Alg. 1 can.

J.24 Testing Membership Inference Attacks

In Tab. J.27, we test membership inference attacks (Shokri et al., 2017) on models finetuned with Alg. 1, i.e., models with uncertainty induced on the protected set. As discussed in Appx. E, we find that Alg. 1 introduces a trade-off with membership inference. Specifically, we train a classifier to map features of the softmax probabilities (the probabilities, their entropy, whether the instance is classified by the model correctly, etc.) to whether or not the instance is in the protected set. We find that this classifier, on a held-out set of 20 examples, is better than random at classifying instances as in the protected set.

J.25 Testing A Finetuning TTP Attack

In Tab. J.28, we test a TTP attack based on finetuning. Specifically, we finetune a model that has induced uncertainty on the protected set with Alg. 1 with a 100 (MNIST) or 300 (IFAR100) test instances for 10 epochs to see if the uncertainty becomes undone. We find that the uncertainty only minimally becomes undone for MNIST, on both Alg. 1 and Alg. 2. For CIFAR100, while the confidence distance increases, the retain and test accuracy drops significantly, suggesting that induced uncertainty is hard to undo by finetuning.

Dataset	Model	MIA Acc.
CIFAR100	ResNet50	65.0%
MNIST	ResNet18	65.0%

Table J.27: We test a membership inference attack on the protected set on models finetuned with Alg. 1, training a classifier to map features of the softmax probabilities to whether or not an instance is in the protected set. As discussed in Appx. E, we find that Alg. 1 introduces a trade-off with membership inference.

Dataset	Model	Method	Stage	Retain Acc.	Test Acc.	Conf. Dist. (Lower Better)
MNIST	LogReg	Pretrain	—	92.1%	91.8%	0.807
		Alg. 1	Initial	87.8%	87.7%	0.180
			After 100-test finetune	88.6%	89.0%	0.350
		Alg. 2	Initial	87.1%	87.2%	0.280
			Attacked	88.4%	88.8%	0.320
CIFAR100	ResNet50	Pretrain	—	100.0%	78.3%	0.902
		Alg. 1	Initial	91.4%	65.5%	0.298
			Attacked	14.8%	14.2%	0.346

Table J.28: We test a TTP attack based on finetuning across dataset and model pairs. In the case of MNIST, confidence distance only increases marginally. In the case of CIFAR100, confidence distance increases at significant cost to retain and test accuracy.

K Symbol Table

Symbols	
f	A pretrained classifier.
f_u	A pretrained classifier after unlearning has been conducted over $\boldsymbol{x}_p$.
$\boldsymbol{x}_p$	A data instance corresponding to person p.
$\mathcal{X}$	A sample space, subset of $\mathbb{R}^d$.
$\mathcal{Y}$	A label space, subset of $\mathbb{R}^o$.
$\mathcal{Z}$	The Cartesian product of a sample space and a label space. This is the space where a dataset is drawn from.
$\mathcal{D}$	A dataset, subset of $\mathcal{Z}^n$, which is the n-fold Cartesian product of $\mathcal{Z}$. This represents a set of n data instances.
$\mathcal{D}_p$	A protected set, a subset of a dataset $\mathcal{D}$.
$\mathcal{D}_r$	A retain set, the complement of the protected set in $\mathcal{D}$.
$\mathcal{W}$	A space of parameters, subset of $\mathbb{R}^p$.
$\mathcal{A}$	A function that maps datasets to parameters; this represents the learning algorithm that a ML model provider uses throughout our paper.
$\mathcal{H}_{\mathcal{W}}$	A set of functions which map samples in $\mathcal{X}$ to the probability simplex $\Delta_{\|\mathcal{Y}\|}$, parameterized by a $\boldsymbol{w} \in \mathcal{W}$.
$\|\boldsymbol{w}\|_2$	The ℓ_2 norm of a vector $\boldsymbol{w}$.
$\|A\|_2$	The 2-operator norm of a matrix $\boldsymbol{A}$.
λ	An ℓ_2 regularization coefficient used in Alg. 1 for regularization and the certified algorithms e.g. Alg. 2 for local convex approximation.
$\lambda_{\min}(\boldsymbol{A})$	The minimum eigenvalue of a matrix $\boldsymbol{A}$.
$\mathcal{K}$	A uniform learner, which maps samples to parameters which, when one parametrizes a function by any such parameter, a uniform distribution over all possible labels, $U[0, \|\mathcal{Y}\|]$ is outputted.
$f_{\boldsymbol{w}}$	A classifier parameterized by a parameter $\boldsymbol{w} \in \mathcal{W}$.
$\mathcal{L}_{\mathcal{A}}$	A loss function to yield accurate predictions e.g. the cross entropy loss between model predictions and labels.
$\mathcal{L}_{\mathcal{K}}$	A loss function to yield accurate uniformity e.g. the Kullback-Liebler divergence between softmax outputs and uniform distribution.
θ	A trade off parameter in $(0, 1)$ between utility and uniformity.

Symbols

$\mathcal{M}_\theta$	A map between datasets and parameters that is the minimizer of a Pareto objective between $\mathcal{L}_\mathcal{A}$ and $\mathcal{L}_\mathcal{K}$, where θ spans the (convex) Pareto frontier.
J	The bound on the norm of the Hessian at $\boldsymbol{w}^*$.
$\mathcal{D}_p^{(i)}$	The ith instance of the protected set.
$\mathcal{D}_r^{(j)}$	The jth instance of the retain set.
$\mathcal{D}_r^{(j,\mathcal{X})}$	The feature of the jth instance of the retain set.
$\mathcal{D}_r^{(j,\mathcal{Y})}$	The label of the jth instance of the retain set.
$\approx_{\varepsilon,\delta,\mathcal{T}}$	Used to denote when two algorithms are (ε, δ) indistinguishable across all subsets $\mathcal{T} \subset \mathcal{W}$, i.e. $\mathcal{M}(\mathcal{D}) \approx_{\varepsilon,\delta,\mathcal{T}} \mathcal{M}'(\mathcal{D}')$ means that $\Pr[\mathcal{M}(\mathcal{D}) \in \mathcal{T}] \leq e^\varepsilon \Pr[\mathcal{M}'(\mathcal{D}') \in \mathcal{T}] + \delta$ and $\Pr[\mathcal{M}'(\mathcal{D}') \in \mathcal{T}] \leq e^\varepsilon \Pr[\mathcal{M}(\mathcal{D}) \in \mathcal{T}] + \delta$.
C	A bound on the model weights.
$P_\mathcal{K}$	The Lipschitz constant for the gradients of $\ell_\mathcal{K}$, the component loss functions of $\mathcal{L}_\mathcal{K}$, from Asm. 5.3.
$P_\mathcal{A}$	The Lipschitz constant for the gradients of $\ell_\mathcal{A}$, the component loss functions of $\mathcal{L}_\mathcal{A}$, from Asm. 5.3.
$F_\mathcal{K}$	The Lipschitz constant for the Hessians of $\ell_\mathcal{K}$, the component loss functions of $\mathcal{L}_\mathcal{K}$, from Asm. 5.4.
$F_\mathcal{A}$	The Lipschitz constant for the Hessians of $\ell_\mathcal{A}$, the component loss functions of $\mathcal{L}$, from Asm. 5.4.
P	A convex combination of $P_\mathcal{K}$ and $P_\mathcal{A}$ with respect to θ.
F	A convex combination of $F_\mathcal{K}$ and $F_\mathcal{A}$ with respect to θ.
$\mathcal{N}(0, \sigma^2 \boldsymbol{I})$	The standard normal distribution with an isotropic covariance matrix.
$\nabla_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$	The gradient of the Pareto objective evaluated at $\boldsymbol{w}^*$, used in the main paper with regularization.
$\boldsymbol{H}_{\boldsymbol{w}^*, \mathcal{K}, \mathcal{A}}$	The Hessian of the Pareto objective evaluated at $\boldsymbol{w}^*$, used in the main paper without regularization.